working paper
department
of economics

PERFECT BAYESIAN AND SEQUENTIAL EQUILIBRIA: A CLARIFYING NOTE

By

Drew Fudenberg and Jean Tirole

No. 496 May 1988

massachusetts
institute of
technology
50 memorial drive
cambridge, mass. 02139

PERFECT BAYESIAN AND SEQUENTIAL EQUILIBRIA: A CLARIFYING NOTE

By

Drew Fudenberg and Jean Tirole

No. 496 May 1988

PERFECT BAYESIAN AND SEQUENTIAL EQUILIBRIA: A CLARIFYING NOTE*

By

Drew Fudenberg**

and

Jean Tirole**

May 1988

* Research support from National Science Foundation Grants SES-8609697 and SES-97619 is gratefully acknowledged. We thank David Kreps for prompting us to write this paper.

** Department of Economics, Massachusetts Institute of Technology

ABSTRACT

We introduce a new and more restrictive notion of perfect Bayesian equilibrium (PBE). A PBE is a specification of strategies and beliefs such that (P) at each stage the strategies form a Bayesian equilibrium for the continuation game, given the specified beliefs, and (B) beliefs are updated from period to period in accordance with Bayes rule whenever possible, and also satisfy a "no-signaling-what-you-don't-know" condition that for all players i and j, the conditional distribution over i's type given j's type and the observed history is independent of player j's actions. This condition can be verified without constructing the sequences of beliefs required to show that an equilibrium is sequential. PBE is equivalent to sequential equilibrium in multistage games provided that each player has only two possible types; the concepts differ when the number of types per player is larger.

1. Introduction.

Kreps-Wilson's [1982] notion of consistent beliefs and sequential equilibria provide one answer to the question "what are reasonable inferences for a player to make if he sees an opponent play an action that has zero probability according to the equilibrium strategies?" Roughly, their answer is that following a deviation players infer (1) that all players will continue to follow the equilibrium strategies and (2) that the deviation was the result of a random mistake or "tremble," where the trembles are independent between information sets, and a player's probability of trembling is measurable with respect to his own information. Beliefs are "consistent" if they can be derived using Bayesian inference from arbitrarily small trembles. A combination of a strategy selection and a system of beliefs is a sequential equilibrium if the beliefs are consistent, and if the strategies are "sequentially rational" in the sense that at every information set the player's strategy maximizes his expected payoff given his beliefs and the strategies of his opponents.

In several applications of dynamic games of incomplete information, economists have used the weaker and vaguer equilibrium concept "perfect Bayesian equilibrium" or "PBE" instead of sequential equilibrium. In a PBE, the strategies are required to be sequentially rational given the beliefs, but the restrictions on beliefs are expressed less formally, and in particular without reference to the sequences of trembles required by the definition of consistency. In the weakest version of PBE, no restrictions at all are placed on the beliefs off of the equilibrium path. While this weak version is sometimes appropriate, in other situations economists may prefer a more restrictive version of PBE that is closer to the sequential equilibrium concept. This paper develops such a version.

The idea behind the restrictions we impose is that a player's deviation should not signal information that the player himself does not possess. (More formally, the conditional distribution over player i's type given player j's type and the observed history is independent of player j's choice of actions.)

We show that this more restrictive definition of PBE is equivalent to sequential equilibrium for the class of multistage games of incomplete information provided that each player is one of two possible "types;" sequential equilibrium imposes additional restrictions when the number of types is three or more. We are agnostic about the plausibility of these additional restrictions given that the PBE restrictions are judged to be reasonable. Our work can be viewed either as proposing a new and less restrictive equilibrium concept or as providing further explanation of what the sequential equilbrium concept implies. Intuitively, sequential equilibrium requires that at every period there be a consistent ranking of the probability of each player's zero-probability types, i.e. some of the zero-probability types are "infinitely more likely" than others. This restriction follows from the requirement that beliefs be consistent with the "trembles" explanation of deviations, as in the games with the trembles the zero-probability types θ_i are assigned positive probabilities ϵ_i^n that converge to zero. These ϵ_i^n define the "infinitely more likely" relationship as follows: θ_i is infinitely more likely than θ_i' if $\epsilon_i^{n'}/\epsilon_i^n \to 0$. The fact that the infinitely-more-likely relationship is an ordering has implications for which beliefs are consistent that are not captured by our PBE concept. These additional restrictions may or may not be reasonable, but they seem different in kind and spirit from the desiderata Kreps-Wilson used to motivate the sequential equilibrium concept.

In the games we consider, play takes place in a number of "periods", with each period's play revealed before proceeding to the next. The only asymmetry of information is that, following Harsanyi [1967-68], each player has a "type" that is chosen by nature at the start of play and revealed only to him. While this is certainly very restrictive, it is a broad enough class to include the literatures on bargaining, limit pricing, and predation, to name but three of the topics to which the theory of games of incomplete information have been applied.

Section 2 describes multistage games of incomplete information, and section 3 derives our equivalence result when the types are independent. Section 4 generalizes the model and results to correlated types.

2. Multi-Stage Games of Incomplete Information.

We will consider only a restricted class of games; that of multi-stage games of incomplete information. This class was first analyzed by Fudenberg-Levine [1983]. Players are denoted by $i \in I$. Each player i has a type θ_i which is drawn from a finite set Θ_i. The prior distribution over the types, denoted $\rho(\theta)$, is assumed to be common knowledge. We assume for the moment that the types are independent, so that $\rho = \prod_i \rho_i$, where ρ_i is the marginal distribution over player i's type. At the beginning of the game, each player is told his own type, but is not given any information about the types of his opponents. That is, player i's partition of θ when his type is θ_i is $\mathcal{P}_i(\theta_i) = \{\theta' | \theta'_i = \theta_i\}$. For notational simplicity we identify the set of player i's partitions with the set θ_i of his types. As the types are independent, player i's initial beliefs about the types of his opponents are given by the prior distribution $\rho_{-i} = \prod_{j \neq i} \rho_j(\theta_j)$.

The game is played in "periods", with the property that at each period t, all players simultaneously choose an action which is then revealed at the end of the period. (This specification is more general than it may appear, because the set of feasible actions can be time and history dependent, so that games with alternating moves are included.) We assume that players never receive additional observations of θ. For notational simplicity, we assume that each player's possible actions are independent of his type. Let $h^0 = 0$, and let $A_i(h^0)$ be the set of player i's possible first-period actions. If the history of moves (other than Nature's choice of types) before t is h^{t-1}, then player i's period-t action must belong to $A_i(h^{t-1})$; if $a^t \in \underset{i \in I}{X} A_i(h^{t-1})$ is played at time t, we set $h^t = (h^{t-1}, a^t)$. We assume that the action sets are finite, and that every player always has at least one feasible action. Since each player i knows θ_i but not θ_{-i}, the information set corresponding to player i's move in period t is identified with an element of $H^{t-1} \times \Theta_i$.

A strategy π_i for player i is a sequence of maps $\pi_i(a_i | h^{t-1}, \theta_i)$ from $H^{t-1} \times \Theta_i$ to $\Sigma_i(h^{t-1})$, where $\Sigma_i(h^{t-1})$ is the space of mixed strategies over $A_i(h^{t-1})$. Player i's payoff $u_i(h^T, \theta_i, \theta_{-i})$ depends on the final history h^T, his own type, and the types of his opponents θ_{-i}. (Note that the payoffs need not be separable over periods.) In a Bayesian Nash equilibrium (Harsanyi [1967-68]) each player's strategy maximizes his expected payoff given his opponents' strategies and his prior beliefs about their types. To extend the spirit of subgame perfection to these games, we would like to require that the strategies yield a Bayesian Nash equilibrium, not only for the whole game, but also for the "continuation games" starting in each period t after every possible history h^{t-1}. Of course, these continuation games are not "proper subgames," because they do not stem from a singleton information set. Thus to make the continuation games into true games we must specify the player's prior beliefs at the start of each continuation game. *A priori* it is not obvious that these beliefs must be common knowledge, but in accordance with most work on refinements we will assume that they are. (One exception to this common-knowledge requirement is the notion of a "c-perfect equilibrium," which is developed in Fudenberg-Kreps-Levine [1987].) Under the common knowledge assumption, the beliefs at the beginning of period t can be represented by a single map μ from H^{t-1} to $\Delta(\Theta)$, the space of probability distributions over Θ. We will denote this distribution by $\mu(\theta | h^{t-1})$, and also use the alternative notation $\mu(h^{t-1})$. Since each player i knows his own type, player i's beliefs about the types of his opponents are given by the conditional distribution $\mu(\theta_{-i}|\theta_i, h^{t-1})$. (In the independent case, this conditional probability is simply the product of the marginal distributions over the types of all players but player i.) Later we will require that the strategies following history h^{t-1} should yield a Bayesian equilibrium relative to the beliefs $\mu(\theta|h^{t-1})$ for all histories h^{t-1}, including those which have zero probability according to the equilibrium strategies.

3. Reasonable Beliefs.

Which systems of beliefs are reasonable? A minimal requirement is that beliefs should be those given by Bayes rule where Bayes rule is applicable, i.e. along the equilibrium path. [This weak requirement plus a twist similar to the no-signaling condition defined below corresponds to the definition of perfect Bayesian equilibrium given in our [1983] paper.] However, as Kreps-Wilson point out, it may often be natural to impose further restrictions on the beliefs.

One additional restriction is that at any date t with beliefs μ^t, the beliefs at date (t+1) should be consistent with Bayes rule applied to the given strategies and the beliefs μ^t, even if those strategies assign probability zero to history h^{t-1}. To motivate this restriction, consider a game where player one has two types, $\underline{\theta}$ and $\bar{\theta}$. Fix an equilibrium where no type of player 1 plays a certain action a_1 in the first period. Since Bayes law does not determine player 2's beliefs when this occurs, we can specify that following this deviation player 2 thinks player 1 is type $\underline{\theta}$. Now let the equilibrium strategies predict that if player 1 does play action a_1 in the first period, he will play b_1 in the second period regardless of his type. It might then seem natural that player 2's beliefs at the start of period 3 when player 1 has played a_1 and then b_1 should be the same as his beliefs at the start of period 2, i.e that player 1 is type $\underline{\theta}$. However, since player 2's corresponding information set is off of the equilibrium path, other beliefs for player 2 would not violate Bayes law. For example, we could specify that player 2 is now certain that player 1 is type $\tilde{\theta}$. To rule this out this reversal of beliefs, we need to require that each player's beliefs in each period are consistent with Bayes law applied to his beliefs of the previous period and the equilibrium strategies. This restriction is related to Kreps-Wilson's notion of structural consistency and also their conditions (5.3) and (5.4).

A more subtle condition is that no player i's actions be treated as containing information about things that player i does not know. This condition is obviously satisfied along the equilibrium path; its role is to constrain the out-of-equilibrium inferences. If deviations from equilibrium are thought of as "random errors," the condition corresponds to the assumption that each player's probability of error depends only on factors known to that player. In a multistage game with independent types, the condition is equivalent to the requirements that the beliefs about different players' types are independent and that each player's deviations are taken as signals only about that player's type. (Section 4 treats the case of correlated types, where one player's deviation can signal information about the type of another.)

<u>Definition</u>: An assement (μ,π) is <u>believable</u> if

(1) Bayes' rule is used to update beliefs wherever possible:
For each $a_i \in A_i(h^{t-1})$, if $\exists \theta_i' \in$ support $\mu_i(\theta_i|h^{t-1})$ such that $\pi_i(a_i|\theta_i',h^{t-1}) > 0$, then for every a^t with $a_i^t = a_i$,

$$\mu_i(\theta_i|(h^{t-1},a^{t-1})) =$$

$$\Pi_i(a_i|\theta_i,h^{t-1})\mu_i(\theta_i|h^{t-1}) \Big/ \Big[\sum_{\tilde{\theta}_i \in \Theta_i} \Pi_i(a_i|\tilde{\theta}_i,h^{t-1})\mu_i(\tilde{\theta}_i|h^{t-1}) \Big].$$

(2) The posterior beliefs at each date are that types are independent:

$$\mu(\theta|h^t) = \prod_i \mu_i(\theta_i|h^t).$$

(3) The beliefs $\mu_i(\theta_i|h^t)$ about player i at period t+1 depend only on $\mu_i(\theta_i|h^{t-1})$ and on player i's period-t action a_i^t:

$$\mu_i(\theta_i|(h^{t-1},a^t)) = \mu_i(\theta_i|(h^{t-1},\bar{a}^t)) \text{ if } a_i^t = \bar{a}_i^t.$$

Conditions (2) and (3) combined are the "no-signalling-what-you-don't-know condition."

"Believability" allows the beliefs about player i's type at time t+1 to be <u>completely</u> <u>arbitrary</u> following a move by player i that has zero conditional probability according to $(\mu,\pi)(h^{t-1})$. The only constraints are that player i's actions not change the beliefs about player j's type, and that, after player i's first zero-probability move results in some new beliefs about his type, subsequent beliefs are determined by Bayes rule and the strategy π until player i deviates again. Thus it is easy to check whether an assessment is believable.

<u>Definition</u>: A perfect Bayesian equilibrium (PBE) is an assessment (μ,π) satisfying:

(B) (μ,π) is believable, and

(P) For each period t and history h^{t-1}, the continuation strategies are a Bayesian equilibrium for the continuation game given the beliefs

$\mu(h^{t-1})$.

While believability captures the three criteria we have discussed, those three do not exhaust the implications of sequential equilibria in general games. Kreps and Wilson say that an assessment (μ,π) is <u>consistent</u> if there is a sequence of totally mixed behavior strategies $\pi^n \rightarrow \pi$ such that the beliefs μ^n computed from π^n using Bayes rule converge to μ. We will say that (μ^n,π^n) "justifies" (μ,π). (A strategy is totally mixed if at every information set the associated behavioral strategy puts strictly positive probability on every action. Thus the beliefs associated with a totally mixed strategy are completely determined by Bayes rule. Note that in our context a totally mixed strategy is one in which at each period t for every history h^{t-1} every type of each player i assigns positive probability to every action in $A_i(h^{t-1})$. Remember also that in a totally mixed behavior strategy the randomizations by different players are independent, as are the randomizations of a single player in different periods.)

<u>Proposition 1</u>: Suppose that each player has only two possible types, that both types have nonzero prior probability and that types are independent. Then an assessment (μ,π) is consistent iff it is believable.

<u>Corollary</u>: Under the hypotheses of Proposition 1, the sets of perfect Bayesian equilibria and sequential equilibria coincide.

<u>Remark</u>: The condition that both types have positive prior probability is necessary because PBE would permit a type with zero prior probability to have positive posterior following a zero-probability action. This is not possible in a sequential equilibrium, as Nature's moves are not subject to trembles.

This is related to the fact that the set of sequential equilibria can change discontinuously when a type is added whose prior probability is arbitrarily small (Fudenberg-Kreps-Levine [1987]).

Proof: If (μ,π) is consistent, fix a sequence of totally mixed strategies $\pi^n \to \pi$ with associated beliefs $\mu^n \to \mu$. Since the π^n correspond to independent randomizations, each player's strategy π_i^n depends only on his type θ_i and the public history h^{t-1}, and the types are independent, the μ^n satisfy (1) and (2). Since these properties are preserved in passing to the limit, consistent beliefs are believable.

Conversely, imagine that (μ,π) is believable. We establish the following claim by induction on T, the number of periods:

Claim: In a T-period game with initial beliefs $\mu(\theta|h^0) = \mu(h^0)$, if (μ,π) is believable then for any strictly positive prior assessment $\mu^n(h^0) \to \mu(h^0)$ there exists a sequence of totally mixed strategies $\pi^n \to \pi$ such that the beliefs μ^n computed from $(\mu^n(h^0), \pi^n)$ using Bayes rule converge to the specified beliefs μ at every information set. Moreover, if $\mu(h^0)$ is strictly positive (i.e., has full support), we can take $\mu^n(h^0) = \mu(h^0)$.

Note that proving this claim is sufficient for our result as the prior distribution is assumed to put positive probability on all types in θ. The reason we consider sequences $\mu^n(h^0)$ converging to $\mu(h^0)$ as opposed to simply $\mu(h^0)$ itself is that we will proceed by induction: First we will construct first-period trembles, then second-period trembles, and so on. In this process we will wish to treat h^0 as the initial history in a continuation game, and in so doing we will need to use the beliefs corresponding to trembles in earlier periods.

Proof of Claim:

I. We begin with a 2-period game, where each player i has two possible types $\underline{\theta}_i$ and $\bar{\theta}_i$. Here the only beliefs which are relevant are those following the first period's play h^1, $\mu(\theta|h^1)$. Because μ is believable, the beliefs about player i depend on $h^1 = (a_1^1, a_2^1)$ only through player i's choice of action a_i^1. In the obvious notation, we let $\underline{\mu}_i(a_i^1) = \mu(\underline{\theta}_i|a_i^1)$; we define $\bar{\mu}_i$, $\bar{\pi}_i$, and $\underline{\pi}_i$ analogously.

Choose a sequence $\epsilon^n \to 0$ and let $\mu^n(h^0) \to \mu(h^0)$ be such that $\mu^n(\theta_i|h^0) > \epsilon^n$. Without loss of generality we assume that $\bar{\mu}_i(h^0) > 0$ for all i. For each player i we define the sets $\bar{P}_i$ of actions in the support of $\bar{\pi}_i$, and $\bar{O}_i$ of actions that $\bar{\pi}_i$ assigns zero probability; the sets $\underline{P}_i$ and $\underline{O}_i$ are defined analogously. We will now construct totally mixed strategies $\pi^n \to \pi$ such that the associated posterior beliefs $\mu^n(h^1)$ computed from $\mu^n(h^0)$ and π^n using Bayes rule converge to the specified posterior μ. To do this we construct π_i^n separately for each player i, beginning with those pairs (a_i^1, θ_i) for which $\pi^i(a_i^1, \theta_i) = 0$. That is, we first construct the "trembles" for type $\underline{\theta}_i$ and actions in $\underline{O}_i$ and type $\bar{\theta}_i$ and actions in $\bar{O}_i$; the strategies π_i^n for other action-type pairs are constructed by substracting the "trembles" assigned to the "zero-probability" actions.

(a) Let us specify the probabilities that player i uses action $a_i^1 \in \underline{O}_i \cap \bar{O}_i$. If $\bar{\mu}_i(a_i^1)$ is positive, we choose $\bar{\pi}_i^n(a_i^1) \to 0$ and $\underline{\pi}_i^n(a_i^1) \to 0$ so that

(i) $\underline{\mu}_i(a_i^1) / \bar{\mu}_i(a_i^1) = \lim_{n\to\infty} \underline{\mu}_i^n(h^0)\ \underline{\pi}_i^n(a_i^1) / \bar{\mu}_i^n(h^0)\ \bar{\pi}_i^n(a_i^1)$.

If $\bar{\mu}_i(a_i^1) = 0$, we choose the π_i^n so that

(ii) $\bar{\mu}_i(a_i^1)/\underline{\mu}_i(a_i^1) = \lim_{n\to\infty} \bar{\mu}_i^n(h^0)\bar{\pi}_i^n(a_i^1)/\underline{\mu}_i^n(h^0)\ \underline{\pi}_i^n(a_i^1)$.

(b) Now we consider actions $a_i^1 \in \underline{P}_i \cap \bar{O}_i$ and construct $\bar{\pi}_i^n$ but do not yet specify π_{-i}^{n}. If $\mu_{-i}(h^0) > 0$, then (μ,π) believable implies that $\mu_{-i}(a_i^1) = 1$. Let $\bar{\pi}_i^n(a_i^1) \to 0$, and note for future reference that as long as $\pi_{-i}^n(a_i^1) \to \pi_{-i}(a_i^1)$ the beliefs μ^n corresponding to action a_i^1 will converge to a point mass on θ_{-i} as desired. If $\mu_{-i}(h^0) = 0$, then we have two cases depending on whether $\bar{\mu}_i(a_i^1)$ is nonzero or zero. If it is nonzero, we choose $\bar{\pi}_i^n(a_i^1) \to 0$ so that

$$\text{(iii)} \qquad \mu_{-i}(a_i^1) \,/\, \bar{\mu}_i(a_i^1) = \lim_{n\to\infty} \mu_{-i}^n(h^0)\pi_{-i}^n(a_i^1) \,/\, \bar{\mu}_i^n(h^0)\bar{\pi}_i^n(a_i^1).$$

If $\bar{\mu}_i(a_i^1) = 0$, we choose the $\bar{\pi}_i^n$ so that

$$\text{(iv)} \qquad \bar{\mu}(a_i^1)/\mu_{-i}(a_i^1) = \lim_{n\to\infty} \bar{\mu}_i^n(h^0)\bar{\pi}_i^n(a_i^1) \,/\, \mu_{-i}^n(h^0)\pi_{-i}^n(a_i^1).$$

Note that as long as $\pi_{-i}^n \to \pi_{-i}$ equations (iii) and (iv) guarantee that

$$\mu_i^n(a_i^1) \to \mu_i(a_i^1).$$

For actions a_i^1 in $\bar{P}_i \cap O_{-i}$ we know that $\bar{\mu}_i(a_i^1) = 1$. For these actions let π_{-i}^n be any sequence converging to zero; then as long as $\bar{\pi}_i^n \to \bar{\pi}_i > 0$ the posteriors will have the appropriate limit.

Finally we specify $\pi_i^n(a_i^1|\theta_i)$ for actions in $P_i(\theta_i)$.

$$\pi_i^n(a_i^1|\theta_i) \equiv \pi_i(a_i^1|\theta_i) - \sum_{\hat{a}_i^1 \in O_i(\theta_i)} \pi_i^n(\hat{a}_i^1,\theta_i) \Big/ \#P_i(\theta_i),$$

where $\#P_i(\theta_i)$ denotes the number of actions in $P_i(\theta_i)$. That is, we subtract the trembles on "zero-probability actions" from the positive probabilities. Note that, for all θ_i,

$$\sum_{a_i^1 \in A_i^1} \pi_i^n(a_i^1|\theta_i) = \sum_{a_i^1 \in A_i^1} \pi_i(a_i^1,\theta_i) = 1.$$

Since $\pi_i^n(a_i^1, \theta_i) \to 0$ for all $a_i \in O_i(\theta_i)$, $\pi_i^n \to \pi_i$. By construction, the posteriors $\mu^n(h^1)$ obtained by updating $\mu^n(h^0)$ using π^n converge to μ. Finally, note that if $\mu(h^0)$ assigns positive probability to both types of θ_i, we can take $\mu^n(h^0) = \mu(h^0)$. This proves our claim for two-period games.

II. Now we extend the claim to games of arbitrary finite length by induction. Assume that the claim is true for all games with T or fewer periods, and consider a game G with T+1 periods along with a believable assesment (μ,π). Let G_{T-1} be the game from period 1 through period (T-1). By inductive hypotheses, there exist initial beliefs $\mu^n(h^0)$ and totally mixed strategies π^n_{T-1} of G_{T-1} such that the associated posterior beliefs μ^n converge to μ at every information set through period T-1. Given the beliefs $\mu^n(h^{T-1})$ at the start of period T, we must show how to choose the period-T strategies $\pi^n_i(a^T_i|h^{T-1},\theta_i)$ so that the posterior beliefs at the start of the period T+1 converge to $\mu(h^T)$ for each $h^T = (h^{T-1},a^T)$. If we then specify that players follow π^n_{T-1} in through period T-1 and $\pi^n_i(a^T_i|h^{T-1},\theta_i)$ at period T, we will have constructed strategies π^n for the T-period game G such that the beliefs computed using prior $\mu^n(h^0)$ and strategies π^n converge to μ at every information set.

Since the beliefs $\mu^n(h^0)$ in part I were an arbitrary sequence of strictly positive vectors converging to $\mu(h^0)$, we can construct the period-T trembles exactly as above.

4. Correlated Types.

When the players' types are not drawn from independent distributions, player i's actions in general signal not only his type, but also those of players whose types are correlated with his. We now generalize the definition of PBE and the equivalence result of section 3 to this situation.

As before, let $\mu(\theta|h^{t-1})$ denote the joint probability distribution at the beginning of period t. For each two-element partition (θ^Y,θ^Z) of θ, we will let $\mu(\theta^Y|\theta^Z,h^{t-1})$ denote the <u>conditional</u> probability of θ^Y given θ^Z. We will also use the <u>marginal</u> distribution $\mu(\theta_i|h^{t-1})$ over a given player's type.

Our generalization of the no-signaling-what-you-don't-know condition is that the beliefs about a set of players at the beginning of period t conditional on the types of the other players depend only on (i) the beliefs at the beginning of period t-1 and (ii) the date (t-1) strategies and date (t-1) actions of this subset of players. For instance, when one player deviates from his equilibrium strategy, he does not affect the probability distribution on the other players' types, conditional on his type. More precisely, an assessment (μ,π) is *believable* if for all subsets of types (θ^Y,θ^Z):

(1′) (Bayes rule): If there exists a vector of types $\hat{\theta}^Y$ such that

$$\mu(\hat{\theta}^Y|\theta^Z,h^{t-1}) > 0 \text{ and } \pi^Y(a^Y|\hat{\theta}^Y,h^{t-1}) > 0$$

(that is, when the history is h^{t-1}, there is positive probability that the types $\hat{\theta}^Y$ will play a^Y at date t.)

$$\mu(\theta^Y|\theta^Z,(a^Y,h^{t-1})) = \frac{\mu(\theta^Y|\theta^Z,h^{t-1})\pi^Y(a^Y|\theta^Y,h^{t-1})}{\sum_{\tilde{\theta}^Y} \mu(\tilde{\theta}^Y|\theta^Z,h^{t-1})\pi^Y(a^Y|\tilde{\theta}^Y,h^{t-1})}$$

(2′) $\mu(\theta^Y|\theta^Z,h^t)$ depends only on $\mu(\theta^Y|\theta^Z,h^{t-1})$, $a^{Y,t}$ and $\pi^{Y,t}$, where $a^{Y,t}$ and $\pi^{Y,t}$ refer to the actions and strategies of the players in subset Y.

A *perfect Bayesian equilibrium* is an assessment (μ,π) that is believable, and such that the continuation strategies from any date t and history h^{t-1} on form a Bayesian equilibrium relative to beliefs $\mu(h^{t-1})$.

Note that equations 1′ and 2′ simplify to equations 1 and 2 respectively when the types are independent.

Proposition 2: With two types per player, the sets of perfect Bayesian and sequential equilibrium strategies coincide.

Remark: Note that this is weaker than the assertion that all believable assessments are consistent. The difference is that we will not show that a believable assessment can be justified by trembles at information sets where two or more players have deviated simultaneously.

Proof of Proposition 2:

First note that a consistent assessment is necessarily believable. The point is that because of the independence of trembles, a player in subset Z can only signal something about a player in subset Y only to the extent that his action signals something about himself. A player's signaling about another player is indirect rather than direct.

To prove the converse, that a perfect Bayesian equilibrium strategy is sequential, we note that the set of equilibrium paths of sequential equilibria is not enlarged if we weaken the consistency requirement to allow the posterior beliefs to be arbitrary at all information sets reached after two or more players deviate simultaneously. This is because in testing for sequential rationality, we ask only if any single player would gain by deviating if his opponents all follow the equilibrium strategy. Thus, an assessment (μ,π) which is sequentially rational, and for which there exist independent trembles on the strategies that justify the beliefs at information sets associated with histories involving at most one player's deviating from his equilibrium strategy per period has the same path as a sequential equilibrium. In the following, we show that a perfect Bayesian equilibrium is sequential in this weaker sense, and therefore its strategies are part of a sequential equilibrium.

To do so, we need to define trembles that justify the specified beliefs, which we do by using the marginals $\mu(\theta_i|h^{t-1})$ to construct trembles as in Proposition 1. Suppose that we have proved inductively that those trembles justify the beliefs through period (T-1).

Consider a vector of actions a by all players which has positive probability of being played at (T-1) given the equilibrium strategies and beliefs. Using Bayes rule and the period (T-1) trembles, we can compute the beliefs at the start of period T as follows.

$$\mu^n(\theta|(h^{T-2},a)) = \frac{\mu^n(\theta|h^{T-2})\prod_i \pi_i^n(a_i|\theta_i,h^{T-2})}{\sum_{\tilde{\theta}} \mu^n(\tilde{\theta}|h^{T-2})\prod_i \pi_i^n(a_i|\tilde{\theta}_i,h^{T-2})} .$$

Now $\mu^n(\theta|h^{T-2})$ converges to $\mu(\theta|h^{T-2})$ by the induction hypothesis, and

$$\pi_i^n(a_i|\theta_i,h^{T-2}) \to \pi_i(a_i|\theta_i,h^{T-2})$$

from the construction of the trembles. (From now on we will suppress the dependence of the strategies on the history h^{T-2} to simplify notation.) Finally, by assumption, we are considering an action a which has positive conditional probability given the history, so that the numerator of the expression is strictly positive. Thus $\mu^n(\theta|(h^{T-2},a))$ converges to $\mu(\theta|(h^{T-2},a))$ as n goes to infinity.

Next, consider an $a = (a^{-i},a_i)$ where only player i deviates from his equilibrium strategy. That is, given the history h^{T-2}, a_i has zero probability but the vector of actions chosen by the other players a^{-i} has positive probability. Let $\pi_{-i}^n(a^{-i}|\theta^{-i}) \equiv \prod_{j\neq i} \pi_j^n(a_j|\theta_j,h^{T-2})$. Then

$$\mu(\theta|(h^{T-2},a)) = \frac{\mu^n(\theta|h^{T-2})\pi^n_{-i}(a^{-i}|\theta^{-i})\pi^n_i(a_i|\theta_i,h^{T-2}_i)}{\sum_{\tilde{\theta}} \mu^n(\tilde{\theta})\pi^n_{-i}(a^{-i}|\tilde{\theta}^{-i})\pi^n_i(a_i|\tilde{\theta}_i,h^{T-2})}$$

or

(3)

$$\mu^n(\theta|(h^{T-2},a)) = \frac{\left[\mu^n(\theta^{-i}|\theta_i,h^{T-2})\pi^n_{-i}(a^{-i}|\theta^{-i})\right]\left[\dfrac{\mu^n(\theta_i|h^{T-2})\pi^n_i(a_i|\theta_i,h^{T-2})}{K^n}\right]}{\sum_{\tilde{\theta}}\left[\mu^n(\tilde{\theta}^{-i}|\tilde{\theta}_i,h^{T-2})\pi^n_{-i}(a^{-i}|\tilde{\theta}^{-i})\right]\left[\dfrac{\mu^n(\tilde{\theta}_i|h^{T-2})\pi^n_i(a_i|\tilde{\theta}_i,h^{T-2})}{K^n}\right]}$$

where $K^n = \max_{\tilde{\theta}_i} \{\mu^n(\tilde{\theta}_i|h^{T-2})\pi^n_i(a_i|\tilde{\theta}_i,h^{T-2})\}$. From our assumption that a_i has zero probability, K^n converges to 0 as n tends to infinity. However, each of the sequences $\dfrac{\mu^n(\tilde{\theta}_i|h^{T-2})\pi^n_i(a_i|\tilde{\theta}_i,h^{T-2})}{K^n}$ lies in the compact set [0,1], so we can extract convergent subsequences that we will also denote by n. Note that at least one of these subsequences converges to a strictly positive number, so that the denominator of (3) converges to a strictly positive limit. Furthermore, the ratio

$$\frac{\mu^n(\theta_i|h^{T-2})\pi^n_i(a_i|\theta_i,h^{T-2})}{[\sum_{\tilde{\theta}_i} \mu^n(\tilde{\theta}_i|h^{T-2})\pi^n_i(a_i|\tilde{\theta}_i,h^{T-2})]}$$

converges to $\mu(\theta_i|(h^{T-2},a))$ by the construction of the trembles. This implies that the right-hand side of equation (3) converges to $\mu(\theta|(h^{T-2},a))$, which was to be proven. Q.E.D.

5. Concluding Remarks.

We conclude with some examples to show the differences between PBE and sequential equilibrium. The fragment of a game in Figure 1 depicts a situation where initially player 1 had three possible types, θ_1', θ_1'', and θ_1^*, but where at time t Bayesian inference given the previous play has led to the conclusion that player 1 must be type θ_1^*. The equilibrium strategies at this point, which are written in curved brackets in the figure, are for type θ_1' to play a_1', type θ_1'' to play a_1'' , and θ_1^* to play a_1^*. Since the first two types have zero probability, player two expects to see player 1 play a_1^*. What should he believe if he sees one of the other two actions? The beliefs in the Figure (given in the square brackets) are that if player 2 sees a_1' he concludes he's facing type θ_1'', while a_1'' is taken as a signal that player 1 is type θ_1'. Since our definition of PBE places no constraints on the beliefs about a player who has just deviated, the situation in Figure 1 is consistent with PBE.

However, the situation of Figure 1 cannot be part of a sequential equilibrium. To see this, imagine that there were trembles π^n that converged to the given strategies π and such that the associated beliefs μ^n converged to the given beliefs μ. Let the probability that μ^n assigns to type θ_1' at the period t be ϵ_n' and let the probability of type θ_1'' be ϵ_n''. Since μ^n converges to μ, both ϵ_n' and ϵ_n'' converge to 0, and $\pi^n(a_1'|\theta_1'')$ and $\pi^n(a_1'|\theta_1^*)$ converge to zero as well. Since $\mu^n(\theta_1''|a_1') = \mu^n(\theta_1'')\ \pi^n(a_1'|\theta_1'') \Big/ \sum_{\theta_1} \mu^n(\theta_1)\ \pi^n(a_1',\theta_1)$, in order to have $\mu^n(\theta_1''|a_1')$ converge to 1 it must be that ϵ_n'/ϵ_n'' converges to zero: In order for the beliefs following a_1' to be concentrated on type θ_1'' when θ_1' plays the action with probability one while θ_1'' assigns it probability zero, the prior beliefs must be that θ_1'' is infinitely more likely than θ_1'. On its own this requirement is compatible with sequential equilibrium. However, considering the beliefs following a_1'' leads to the conclusion that ϵ_n''/ϵ_n'

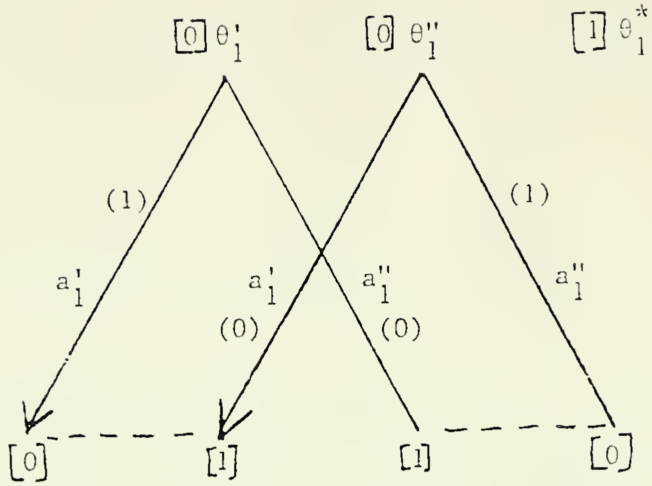

Figure 1

converges to zero, i.e. that θ_1' is infinitely more likely than θ_1'', and these two conditions are jointly incompatible with the beliefs being consistent.

We believe that in games with only three types per payer and only three periods, the following additional condition is sufficient for the equivalence of PBE and sequential equilibrium: For every player i, every pair of types θ_i' and θ_i'', and pair of second-period actions a_i, $\hat{a}_i$,

$$\pi_i(a_i|\theta_i')\mu_i(\theta_i'|a_i)\pi_i(\hat{a}_i,\theta_i'')\mu_i(\theta_i''|\hat{a}_i) = \pi_i(\hat{a}_i,\theta_i')\mu_i(\theta_i'|\hat{a}_i)\pi_i(a_i,\theta_i'')\mu_i(\theta_i''|a_i).$$

This condition ensures that the relationship "infinitely more likely than" which is implicit in the given beliefs and strategies can be extended to an ordering of the relative probabilities of the three types. Note that it is not satisfied in Figure 1. We have not pursued this line of research because in games with more types or periods more complex conditions are needed.

We have described a notion of PBE for multistage games that imposes many of the constraints that Kreps and Wilson argued were desirable. In more complex games, it is harder to impose these constraints without using the apparatus of trembles. Figure 2 depicts a game created by David Kreps in which player 1 moves first, and neither player 2 or player 3 knows which of them is moving second and which is moving third. Our PBE concept cannot be appealed to this game, because its information sets cannot be ordered by precedence, so we can't update the beliefs period-by-period. Sequential equilibrium imposes constraints on the relationships between beliefs at different information sets that seem hard to capture with an extended notion of PBE: The product $\mu_2(a)\mu_2(b)$ equals $\mu_3(c)\mu_3(d)$ (which, since $\mu_2(a) + \mu_2(b) = \mu_3(c) + \mu_3(d) = 1$, implies $\mu_2(a) = \mu_3(c)$ and $\mu_2(b) = \mu_3(d)$).

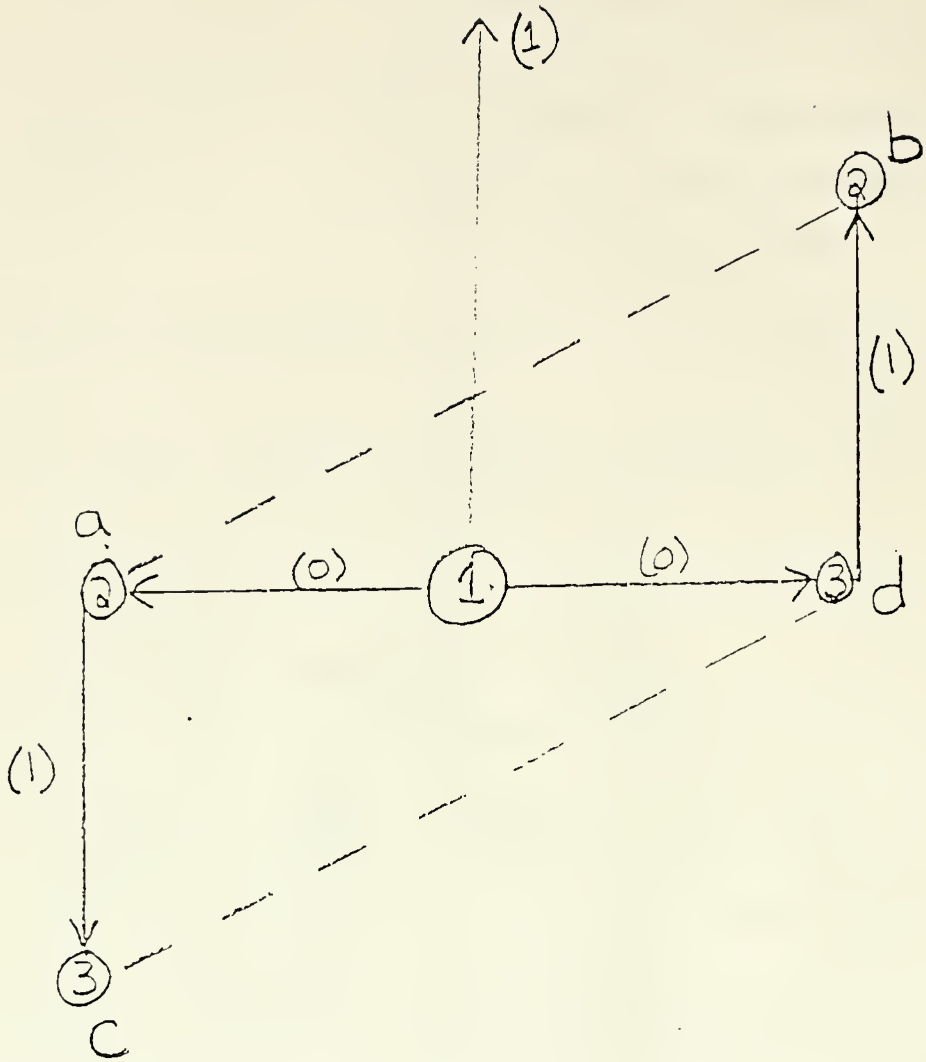

Figure 2

(Dotted lines represent information sets, numbers adjacent to arrows the probability of play. All the actions of players 2 and 3 are not represented).

REFERENCES

Fudenberg, D. and D. Levine, "Subgame-Perfect Equilibria of Finite and Infinite Horizon Games", Journal of Economic Theory, 31 (1983), 251-268.

Fudenberg, D., Kreps, D. and D. Levine, "On the Robustness of Equilibrium Refinements", forthcoming in Journal of Economic Theory.

Fudenberg, D. and J. Tirole, "Sequential Bargaining with Incomplete Information", Review of Economic Studies, 50 (1983), 221-247.

Kreps, D. and R. Wilson, "Sequential Equilibria", Econometrica 50 (1982), 863-894.

3-22-89