MANAGEMENT SCIENCE
Vol. 14, No. 5, January, 1968
Printed in U.S.A.

GAMES WITH INCOMPLETE INFORMATION PLAYED BY "BAYESIAN" PLAYERS

Part II. Bayesian Equilibrium Points*†[1]

JOHN C. HARSANYI

Univerity of California, Berkeley

Part I of this paper has described a new theory for the analysis of games with incomplete information. It has been shown that, if the various players' subjective probability distributions satisfy a certain mutual-consistency requirement, then any given game with incomplete information will be equivalent to a certain game with complete information, called the "Bayes-equivalent" of the original game, or briefly a "Bayesian game."

Part II of the paper will now show that any Nash equilibrium point of this Bayesian game yields a "Bayesian equilibrium point" for the original game and conversely. This result will then be illustrated by numerical examples, representing two-person zero-sum games with incomplete information. We shall also show how our theory enables us to analyze the problem of exploiting the opponent's erroneous beliefs.

However, apart from its indubitable usefulness in locating Bayesian equilibrium points, we shall show it on a numerical example (the Bayes-equivalent of a two-person cooperative game) that the normal form of a Bayesian game is in many cases a highly unsatisfactory representation of the game situation and has to be replaced by other representations (e.g., by the semi-normal form). We shall argue that this rather unexpected result is due to the fact that Bayesian games must be interpreted as games with "delayed commitment" whereas the normal-form representation always envisages a game with "immediate commitment."

8.

Let G be an I-game (as assessed by player j), and let G^* be a Bayesian game Bayes-equivalent to G and having the function $R^*(c_1, \cdots, c_n)$ as its basic probability distribution. Let s_i^* be a normalized strategy of player i in game G. Suppose that, for some specific value $c_i = c_i^0$ of player i's attribute vector c_i, the (ordinary) strategy $s_i = s_i^*(c_i)$ selected by this normalized strategy s_i^* maximizes player i's conditional payoff expectation

$$\mathcal{E}(x_i \mid c_i^0) = Z_i(s_1^*, \cdots, s_i^*, \cdots, s_n^* \mid c_i^0) \tag{8.1}$$

if the normalized strategies $s_1^*, \cdots, s_{i-1}^*, s_{i+1}^*, \cdots, s_n^*$ of the other $(n-1)$ players are kept constant. Then we shall say that this normalized strategy s_i^* of player i is a *best reply at the point* $c_i = c_i^0$ to the other players' normalized strategies $s_1^*, \cdots, s_{i-1}^*, s_{i+1}^*, \cdots, s_n^*$.

* Received June 1965 and revised May 1966.

† Part I of "Games with Incomplete Information Played by 'Bayesian' Players" appeared in the preceding issue of *Management Science*. Part III will appear in the next issue. A "Table of Contents" and "Glossary" appeared with Part I.

[1] The numbering of sections and theorems will be consecutive in Parts I, II, and III of this paper. The author wishes again to express his indebtedness to the persons and institutions listed in Footnote 1 of Part I.

Now suppose that s_i^* actually possesses this best-reply property at *all* possible values of the attribute vector c_i, with the possible exception of a small set C_i^* of c_i values, having a total probability mass zero. That is, we are assuming that the event $E = \{c_i \in C_i^*\}$ is being assigned zero probability by the marginal probability distribution

$$R^*(c_i) = \int_{C^i} d_{(c^i)} R^*(c_i, c^i) \tag{8.2}$$

derived from the basic probability distribution $R^*(c_1, \cdots, c_i, \cdots, c_n) = R^*(c_i, c^i)$. Then we shall say that s_i^* is *almost uniformly* a best reply to $s_1^*, \cdots, s_{i-1}^*, s_{i+1}^*, \cdots, s_n^*$.

Finally suppose that in a given normalized strategy n-tuple $s^* = (s_1^*, \cdots, s_n^*)$ *every* component s_i^* $(i = 1, \cdots, n)$ is almost uniformly a best reply to the other $(n - 1)$ components $s_1^*, \cdots, s_{i-1}^*, s_{i+1}^*, \cdots, s_n^*$. Then we shall say that s^* is a *Bayesian equilibrium point in game G.*

We can now state the following theorem.

Theorem I. Let G be an I-game, and let G^* be a Bayesian game Bayes-equivalent to G (as assessed by player j). In order that any given n-tuple of normalized strategies $s^* = (s_1^*, \cdots, s_n^*)$ be a Bayesian equilibrium point in game G, it is both *sufficient* and *necessary* that in the normal form $\mathfrak{N}(G^*)$ of game G^* this n-tuple s^* be an equilibrium point in Nash's sense [3, 5].

Proof. In view of equations (5.3), (7.4) and (7.7) we can write

$$W_i(s_1^*, \cdots, s_i^*, \cdots, s_n^*) = \int_{C_i} Z_i(s_1^*, \cdots, s_i^*, \cdots, s_n^* \mid c_i)\, d_{(c_i)} R^*(c_i) \tag{8.3}$$

where $R^*(c_i)$ is again the marginal probability distribution defined by equation (8.2).

To prove the sufficiency part of the theorem, suppose that s^* is not a Bayesian equilibrium point in G. This means that at least one of the components of s^*, say, s_i^*, is not an almost uniformly best reply to the other components of s^*. Consequently, there must be some set C_i^* of possible c_i values at which the function $Z_i(\cdot \mid c_i)$ could be increased if we replaced the ordinary strategies $s_i = s_i^*(c_i)$ selected by the normalized strategy s_i^*, with some alternative ordinary strategies $s'_i = s_i^{**}(c_i)$. That is, for all $c_i \in C_i^*$ we must have

$$Z_i[s_i^{**}(c_i), (s^*)^i] > Z_i[s_i^*(c_i), (s^*)^i], \tag{8.4}$$

where $(s^*)^i$ denotes the $(n - 1)$-tuple

$$(s^*)^i = (s_1^*, \cdots, s_{i-1}^*, s_{i+1}^*, \cdots, s_n^*). \tag{8.5}$$

Moreover, the probability distribution $R^*(c_i)$ must associate a non-zero total probability mass with this set C_i^*.

Now let s_i^{***} be a normalized strategy of player i, selecting the strategies

$s_i^{**}(c_i)$ for all points c_i in set C_i^*, but coinciding with s_i^* everywhere else. That is,

$$(8.6) \qquad \begin{aligned} s_i^{***}(c_i) &= s_i^{**}(c_i) \qquad \text{for all } c_i \in C_i^* \\ s_i^{***}(c_i) &= s_i^*(c_i) \qquad \text{for all } c_i \notin C_i^*. \end{aligned}$$

Then, in view of equations (8.3) and (8.4), we must have

$$(8.7) \qquad W_i[s_i^{***}, (s^*)^{i}] > W_i[s_i^*, (s^*)^{i}].$$

Consequently, in game $\mathfrak{N}(G^*)$ the n-tuple s^* is not an equilibrium point in Nash's sense. Thus, if s^* is not a Bayesian equilibrium point in G, then it cannot be a Nashian equilibrium point in $\mathfrak{N}(G^*)$, contrary to our hypothesis.

To prove the necessity part of the theorem, suppose that s^* is not an equilibrium point in Nash's sense in game $\mathfrak{N}(G^*)$. This means that at least one component of s^*, say, s_i^*, can be replaced by some alternative normalized strategy s_i^{***} in such a way as to increase the numerical value of the function W_i. That is, there must exist some normalized strategy s_i^{***} satisfying (8.4). But in view of equation (8.3) this implies that, for some set C_i^* of possible c_i values, we must have

$$(8.8) \qquad Z_i[s_i^{***}(c_i), (s^*)^{i}] > Z_i[s_i^*(c_i), (s^*)^{i}].$$

Moreover, this set C_i^* must have a non-zero total probability mass associated with it under the probability distribution $R^*(c_i)$. Consequently, s^* will not be a Bayesian equilibrium point in game G. Thus, if s^* is not a Nashian equilibrium point in $\mathfrak{N}(G^*)$ then it cannot be a Bayesian equilibrium point in game G, contrary to our hypothesis. This completes the proof.

In view of Nash's [3, 5] equilibrium-point theorem and its various generalizations, any I-game satisfying some rather mild regularity conditions will have at least one Bayesian equilibrium point. In particular:

Theorem II. Let G be a standard-form I-game (as assessed by player j), for which a Bayes-equivalent Bayesian game G^* exists. Suppose that G is a finite game, in which each player i has only a finite number of (ordinary) pure strategies s_i. Then G will have at least one Bayesian equilibrium point.

Proof. Even if each player i has only a finite number of pure strategies s_i in the standard form of game G, he will have an infinity of pure normalized strategies s_i^* in the normal form $\mathfrak{N}(G)$ of the game if his information vector c_i can take an infinity of different values. But it is easy to see that the normalized behavioral strategies of $\mathfrak{N}(G)$ will even then always satisfy the continuity and contractibility requirements of Debreu's generalized equilibrium-point theorem [1]. Hence, $\mathfrak{N}(G)$ will always contain an equilibrium point s^* in Nash's sense, and so by Theorem I s^* will be a Bayesian equilibrium point for game G itself as originally given in standard form.

9.

Now we propose to discuss two numerical examples, partly to illustrate the concept of Bayesian equilibrium points as defined in the last section, and partly to illustrate the nature of the solution concept our approach yields in the simplest class of I-games, viz., two-person zero-sum games with incomplete information.

Suppose that in a given two-person zero-sum game G player 1 can belong to either of two attribute classes $c_1 = a^1$ and $c_1 = a^2$, where his belonging to class a^1 may be intuitively interpreted that he is in a "weak" position (e.g., in terms of military equipment, or man power, or the economic resources available to him, etc.) whereas his belonging to class a^2 may be interpreted that he is in a "strong" position. Likewise, player 2 can belong to either of two attribute classes $c_2 = b_{\mathrm{I}}$ and $c_2 = b^2$, where class b^1 again may be taken to indicate that he is in a "weak" position whereas class b^2 indicates that he is in a "strong" position. Thus, we have altogether four possible cases in that the two players may belong to classes (a^1, b^1) or (a^1, b^2) or (a^2, b^1) or (a^2, b^2).

In any one of these cases, each player has two pure strategies, called $s_1 = y^1$ and $s_1 = y^2$ in the case of player 1, and called $s_2 = z^1$ and $s_2 = z^2$ in the case of player 2. (In terms of an intuitive interpretation, y^1 and z^1 may represent "more aggressive" strategies while y^2 and z^2 may represent "less aggressive" ones.) These strategies will yield the following payoffs to player 1 in the four possible cases (player 2's payoff is always the negative of player 1's).

TABLE 1

In case (a^2, b^1): Saddle point at (y^1, z^1).

	z^1	z^2
y^1	2	5
y^2	−1	20

TABLE 2

In case (a^1, b^2): Saddle point at (y^2, z^1).

	z^1	z^2
y^1	−24	−36
y^2	0	24

TABLE 3

In case (a^2, b^1): Saddle point at (y^1, z^2).

	z^1	z^2
y^1	28	15
y^2	40	4

TABLE 4

In case (a^2, b^2): Saddle point at (y^1, z^1).

	z^1	z^2
y^1	12	20
y^2	2	13

(For the sake of simplicity, in all four cases we have chosen payoff matrices with saddle points in pure strategies.)

Since both players' attribute vectors c_1 and c_2 can take only a finite number (viz. two) alternative values, all probability distributions of the game can be represented by the corresponding probability-mass functions. In particular, the basic probability distribution $R^*(c_1, c_2)$ can be represented by the matrix of joint probabilities r_{km} = Prob ($c_1 = a^k$ and $c_2 = b^m$) with $k, m = 1, 2$. We shall assume that this matrix $\{r_{km}\}$ will be as follows.

TABLE 5

	$c_2 = b^1$	$c_2 = b^2$
$c_1 = a^1$	$r_{11} = 0.40$	$r_{12} = 0.10$
$c_1 = a^2$	$r_{21} = 0.20$	$r_{22} = 0.30$

This matrix yields the following conditional probabilities, corresponding to the two *rows*:

TABLE 6

	$c_2 = b^1$	$c_2 = b^2$
Prob ($c_2 = b^m \mid c_1 = a^1$)	$p_{11} = 0.8$	$p_{12} = 0.2$
Prob ($c_2 = b^m \mid c_1 = a^2$)	$p_{21} = 0.4$	$p_{22} = 0.6$

Here $p_{km} = r_{km}/(r_{k1} + r_{k2})$. The first row of this matrix states the conditional probabilities (subjective probabilities) that player 1 will assign to the two alternative hypotheses about player 2's attribute class (viz. to $c_2 = b^1$ and to $c_2 = b^2$) when his own attribute class is $c_1 = a^1$. The second row states the probabilities he will assign to these two hypotheses when his own attribute class is $c_1 = a^2$.

In contrast, the conditional probabilities corresponding to the two *columns* of matrix $\{r_{km}\}$ are:

TABLE 7

	Prob ($c_1 = a^k \mid c_2 = b^1$)	Prob ($c_1 = a^k \mid c_2 = b^2$)
$c_1 = a^1$	$q_{11} = 0.67$	$q_{12} = 0.25$
$c_1 = a^2$	$q_{21} = 0.33$	$q_{22} = 0.75$

Here $q_{km} = r_{km}/(r_{1m} + r_{2m})$. The first column of this matrix states the conditional probabilities (subjective probabilities) that player 2 will assign to the two alternative hypotheses about player 1's attribute class (viz. to $c_1 = a^1$ and to $c_1 = a^2$) when his own attribute class is $c_2 = b^1$. The second column states the probabilities he will use when his own attribute class is $c_2 = b^2$.

All three probability matrices stated in Tables 5 to 7 show that in our example the two players assume a positive correlation between their attribute classes, in the sense that it will tend to *increase* the probability a given player will assign to the hypothesis that his opponent is in a strong (or weak) position if he himself is in a strong (or weak) position. (Such a situation may arise if both players' power positions are determined largely by the same environmental influences; or if each player is making continual efforts to keep his own military, industrial, or other kind of strength more or less on a par with the other player's, etc. Of course in other situations there may be a negative correlation, or none at all.)

In this game, player 1's normalized pure strategies will be of the form $s_1^* = y^{nt} = (y^n, y^t)$, where $s_1 = y^n$ ($n = 1, 2$) is the ordinary pure strategy player 1 would use if $c_1 = a^1$, whereas y^t ($t = 1, 2$) is the ordinary pure strategy he would use if $c_1 = a^2$. Likewise, player 2's normalized pure strategies will be of the form $s_2^* = z^{uv} = (z^u, z^v)$, where $s_2 = z^u$ ($u = 1, 2$) is the ordinary pure strategy player 2 would use if $c_2 = b^1$, whereas $s_2 = z^v$ is the ordinary pure strategy he would use if $c_2 = b^2$.

It is easy to verify that the payoff matrix W for the normal form $\mathfrak{N}(G)$ will be as follows.

TABLE 8

	z^{11}	z^{12}	z^{21}	z^{22}
y^{11}	7.6	8.8	6.2	7.4
y^{12}	7.0	9.1	1.0	3.1
y^{21}	8.8	13.6	14.6	19.4
y^{22}	8.2	13.9	9.4	15.1

The entries of this matrix are player 1's *total* (i.e., unconditional) payoff expectations corresponding to alternative pairs of normalized strategies (y^{nt}, z^{uv}). For instance, the entry for the strategy pair (y^{12}, z^{11}) is:

$$0.4 \times 2 + 0.1 \times (-24) + 0.2 \times 40 + 0.3 \times 2 = 7.0, \quad \text{etc.}$$

The only saddle point of this matrix W in pure strategies is at the point (y^{21}, z^{11}). (There is none in mixed strategies.) This means that in this game player 1's optimal strategy will be to use strategy y^2 if his attribute vector takes the value $c_1 = a^1$, and to use strategy y^1 if his attribute vector takes the value $c_1 = a^2$. On the other hand, player 2's optimal strategy will always be z^1, irrespective of the value of his attribute vector c_2.

It is easy to see that, in general, these optimal strategies are quite different from the optimal strategies associated with the four matrices stated in Tables 1 to 4. For instance, if the payoff matrix for case (a^1, b^1) is considered in isolation (see Table 1) then player 1's optimal strategy is y^1, though if we consider the game as a whole then his optimal normalized strategy y^{21} actually requires him to use strategy y^2 whenever his attribute vector has the value $c_1 = a^1$. Likewise, if the payoff matrix for case (a^2, b^1) is considered in isolation (see Table 3), then player 2's optimal strategy is z^2, though his optimal normalized strategy z^{11} for the game as a whole always requires him to use strategy z^1. These facts, of course, are not really surprising: they show only that the players will have to use different strategies if they do not know each other's attribute vectors, and different strategies if they do know them. For example, in case (a^1, b^1) player 1's optimal strategy would be y^1 only if he did know that his opponent's attribute vector is $c_2 = b^1$. But if all his information about player 2's attribute vector c_2 is what he can infer from the fact that his own attribute vector is $c_1 = a^1$, on the basis of the probability matrices of Tables 5 to 7, then his optimal strategy will be y^2 as prescribed by his optimal normalized strategy y^{21}, etc.

The strategy pair (y^{21}, z^{11}), being a saddle point, will also be an equilibrium point in Nash's sense for the normal form $\mathfrak{N}(G)$ of the game. Thus, so long as player 2 is using strategy z^{11}, player 1 will be maximizing his *total* payoff expectation W_1 by choosing strategy y^{21}. Conversely, so long as player 1 is using strategy y^{21}, player 2 will be maximizing his total payoff expectation W_2 by choosing strategy z^{11}.

What is more important, by Theorem I of the previous section, the strategy pair (y^{21}, z^{11}) will also be a *Bayesian equilibrium point* in the standard form of game G (indeed it will be the only one since it is the only saddle point of the normal-form game). Hence, so long as player 2 uses strategy z^{11}, player 1 will maximize, not only his *total* payoff expectation W_1, but also his *conditional* payoff expectation $Z_1(\cdot \mid c_1)$, *given* the value of his own attribute vector c_1. To verify this we first list player 1's conditional payoff expectations when his attribute vector takes the value $c_1 = a^1$:

TABLE 9

	z^{11}	z^{12}	z^{21}	z^{22}
y^1	−3.2	−5.6	−0.8	−3.2
y^2	−0.8	4.0	16.0	20.8

We shall call this matrix $Z_1(\cdot \mid a^1)$. We do find that the highest entry of column z^{11} is in row y^2, corresponding to the *first* component of the normalized strategy $y^{21} = (y^2, y^1)$.

Next we list player 1's conditional payoff expectations when his attribute vector takes the value $c_1 = a^2$:

TABLE 10

	z^{11}	z^{12}	z^{21}	z^{22}
y^1	18.4	23.2	13.2	18.0
y^2	17.2	23.8	2.8	9.4

We shall call this matrix $Z_1(\cdot \mid a^2)$. Now we find that the highest entry of column z^{11} is in row y^1, corresponding to the *second* component of the normalized strategy $y^{21} = (y^2, y^1)$, as desired. Thus in both cases the ordinary strategy representing the appropriate component of the normalized strategy y^{21} maximizes player 1's conditional payoff expectation against the normalized strategy z^{11} of player 2, in accordance with Theorem I.

We leave it to the reader to verify by a similar computation that conversely the ordinary strategy z^1 [which happens to be both the first and the second component of the normalized strategy $z^{11} = (z^1, z^1)$] does maximize player 2's conditional payoff expectation against the normalized strategy y^{21} of player 1, both in case player 2's attribute vector takes the value $c_2 = b^1$ and in case it takes the value $c_2 = b^2$, as required by Theorem I.

In this particular example, each player's optimal strategy not only maximizes his security level in terms of his *total* payoff expectations W_i (which is true by definition), but also maximizes his security level in terms of his *conditional* payoff expectations $Z_i(\cdot \mid c_i)$. For instance, in Table 9, strategy y^2 (which is the ordinary strategy prescribed for him by his optimal normalized strategy y^{21} in the case $c_1 = a^1$) is not only player 1's *best reply* to the normalized strategy z^{11} of his opponent, but is also his *maximin* strategy. Likewise, in Table 10, strategy y^1 is not only player 1's best reply to z^{11} but is also his maximin strategy.

However, we now propose to show by a counter example that this relationship has no general validity for two-person zero-sum games with incomplete information. Consider a game G^0 where each player's attribute vector can again take two possible values, viz. $c_1 = a^1, a^2$ and $c_2 = b^1, b^2$. Suppose the payoff matrices in the resulting four possible cases are as follows:

TABLE 11

In case (a^1, b^1):

	z^1	z^2
y^1	8	−8
y^2	0	−4

Saddle point at (y^2, z^2).

TABLE 12

In case (a^1, b^2):

	z^1	z^2
y^1	−4	8
y^2	0	12

Saddle point at (y^2, z^1).

TABLE 13

In case (a^2, b^1):

	z^1	z^2
y^1	−8	12
y^2	−12	16

Saddle point at (y^1, z^1).

TABLE 14

In case (a^2, b^2):

	z^1	z^2
y^1	4	−4
y^2	0	−8

Saddle point at (y^1, z^2).

Suppose that basic probability matrix of this game G^0 is as follows:

TABLE 15

	$c_2 = b^1$	$c_2 = b^2$
$c_1 = a^1$	$r_{11} = 0.25$	$r_{12} = 0.25$
$c_1 = a^2$	$r_{21} = 0.25$	$r_{22} = 0.25$

The payoff matrix W for the normal form $\mathfrak{N}(G^0)$ of G^0, listing player 1's *total* payoff expectations, will now be:

TABLE 16

	z^{11}	z^{12}	z^{21}	z^{22}
y^{11}	0	1	1	2
y^{12}	−2	−1	1	2
y^{21}	−1	0	3	4
y^{22}	−3	−2	3	4

The only saddle point of this matrix is at (y^{11}, z^{11}). Hence player 1's optimal normalized strategy is $y^{11} = (y^1, y^1)$ whereas player 2's is $z^{11} = (z^1, z^1)$. Player 1's *conditional* payoff expectations $Z_1(\cdot \mid c_1)$ in the case $c_1 = a^1$ will be as follows:

TABLE 17

	z^{11}	z^{12}	z^{21}	z^{22}
y^1	2	8	−6	0
y^2	0	6	−2	4

In accordance with Theorem I, as the table shows, player 1's best reply to the opponent's optimal normalized strategy z^{11} is strategy y^1, the ordinary strategy corresponding to his own optimal normalized strategy y^{11}. But player 1's security level would be maximized by y^2, and not by y^1.

This result has the following implication. Let G be a two-person zero-sum game with incomplete information. Then we may use the von Neumann-Morgenstern solution of the normal form $\mathfrak{N}(G)$ of this game as the solution for the game G itself as originally defined in standard form. But, in general, the only real justification for this procedure will lie in the fact that every pair of optimal normalized strategies s_1^* and s_2^* will represent a Bayesian equilibrium point in G. However, we cannot justify this approach by the maximin and minimax properties of s_1^* and s_2^* because in general these strategies will not have such properties in terms of the players' *conditional* payoff expectations $Z_i(\cdot \mid c_i)$, which are the quantities the players will be really interested in and which they will treat as their true anticipated payoffs from the game.

10.

We now propose to illustrate on a third numerical example how our model enables us to deal with the problem of *exploiting the opponent's erroneous beliefs*. Let G again be a two-person zero-sum game where each player can belong to two different attribute classes, viz. to class $c_1 = a^1$ or to class $c_1 = a^2$, and to class $c_2 = b^1$ or to class $c_2 = b^2$, respectively. We shall assume that in the four possible combinations of attribute classes the payoffs of player 1 will be again the same as stated in Tables 1 to 4 of the previous section. However, the basic probability matrix $\{r_{km}\}$ of the game will now be assumed to be:

TABLE 18

	$c_2 = b^1$	$c_2 = b^2$
$c_1 = a^1$	$r_{11} = 0.01$	$r_{12} = 0.00$
$c_1 = a^2$	$r_{21} = 0.09$	$r_{22} = 0.90$

Accordingly, player 1 will now use the following conditional probabilities (subjective probabilities):

TABLE 19

	$c_2 = b^1$	$c_2 = b^2$
Prob $(c_2 = b^m \mid c_1 = a^1)$	$p_{11} = 1.00$	$p_{12} = 0.00$
Prob $(c_2 = b^m \mid c_1 = a^2)$	$p_{21} = 0.09$	$p_{22} = 0.91$

On the other hand, player 2 will now use the following conditional probabilities (subjective probabilities):

TABLE 20

	$\text{Prob}(c_1 = a^k \mid c_2 = b^1)$	$\text{Prob}(c_1 = a^k \mid c_2 = b^2)$
$c_1 = a^1$	$q_{11} = 0.10$	$q_{12} = 0.00$
$c_1 = a^2$	$q_{21} = 0.90$	$q_{22} = 1.00$

The payoff matrix W for the normal form of the game, stating player 1's *total* payoff expectations, will now be:

TABLE 21

	z^{11}	z^{12}	z^{21}	z^{22}
y^{11}	13.34	20.54	11.20	19.40
y^{12}	5.42	15.32	2.21	12.11
y^{21}	13.31	20.51	12.35	19.55
y^{22}	5.39	15.29	2.36	12.26

The only saddle point of this game is at (y^{21}, z^{21}). Thus player 1's optimal strategy is $y^{21} = (y^2, y^1)$ whereas player 2's optimal strategy is $z^{21} = (z^2, z^1)$.

Now suppose that player 1 actually belongs to attribute class $c_1 = a^1$. Then our result concerning the players' optimal strategies can be interpreted as follows. As $c_1 = a^1$, by Table 19 player 1 will be able to infer that player 2 must belong to attribute class $c_2 = b^1$. Hence, by Table 20 player 1 will be able to conclude that player 2 is assigning a near-unity probability (viz. $q_{21} = .90$) to the *mistaken hypothesis* that player 1 belongs to class $c_1 = a^2$, which would mean that case (a^2, b^1) would apply (since player 2's own attribute class is $c_2 = b^1$). Therefore, player 1 will expect player 2 to choose strategy z^2, which is the strategy the latter would use by Table 3 if he thought that case (a^2, b^1) represented the actual situation.

If case (a^2, b^1) did represent the actual situation, then by Table 3 player 1's best reply to strategy z^2 would be strategy y^1. But under our assumptions player 1 will know that the actual situation is (a^1, b^1), rather than (a^2, b^1). Hence, by Table 1 his actual best reply to z^2 will be y^2, and not y^1. Thus we can say that by choosing strategy y^2 player 1 will be able to *exploit* player 2's *mistaken belief* that player 1 probably belongs to class $c_1 = a^2$.

As Table 1 shows, if player 2 knew the true situation and acted accordingly then player 1 could obtain no more than the payoff $x_1 = 2$. If player 2 did make the mistaken assumption that player 1's attribute vector was $c_1 = a^2$, but if player 1 did not exploit this mistake (i.e., if player 1 used the same strategy y^1 he would use against a better informed opponent) then his payoff would be $x_1 = 5$. Yet if he does exploit player 2's mistake and counters the latter's stragety z^2 by his strategy y^2, then he will obtain the payoff $x_1 = 20$.

11.

The purpose of our last numerical example will be to show that the normal form $\mathfrak{N}(G) = \mathfrak{N}(G^*)$ is often an unsatisfactory representation of the corresponding

I-game G (as originally defined in standard form), and is an unsatisfactory representation even for the corresponding Bayesian game G^*. More particularly, we shall argue that the solutions we would naturally associate with the normal-form game $\mathfrak{N}(G) = \mathfrak{N}(G^*)$ may often give highly counter-intuitive results for G or G^*.

Let G be a two-person bargaining game in the sense of Nash [4] in which the two players have to divide \$100. If they cannot agree on their shares, then both of them will obtain zero payoffs. It will be assumed that both players have linear utility functions for money, so that their dollar payoffs can be taken to be the same as their utility payoffs.

Following Nash, we shall assume that the game will be played as follows. Each player i ($i = 1, 2$) has to state his payoff demand y_i^* without knowing what the other player's payoff demand y_j^* will be. If $y_1^* + y_2^* \leqq 100$, then each player will obtain a payoff $y_i = y_i^*$ equal to his own demand. If $y_1^* + y_2^* > 100$, then both players will obtain zero payoffs $y_1 = y_2 = 0$.

To introduce incomplete information in the game, we shall assume that either player i may have to pay half of his gross payoff y_i to a certain secret organization or he may not; and that he himself will always know in advance whether this is the case or not, but that neither player will know whether the other has to make this payment or not. We shall say that player i's attribute vector takes the value $c_i = a^1$ if he has to make this payment, and takes the value $c_i = a^2$ if this is not the case.

Thus, we can define the net payoffs x_i of both players i as:

$$(11.1) \qquad \begin{aligned} x_i &= \tfrac{1}{2} y_i \quad \text{if} \quad c_i = a^1 \\ x_i &= y_i \quad \text{if} \quad c_i = a^2 \end{aligned} \qquad i = 1, 2.$$

Let X be the *payoff space* of the game, that is, the set of all possible net-payoff vectors $x = (x_1, x_2)$ that the players can achieve in the game.

If the players' attribute vectors are (a^1, a^1), then $X = X^{11}$ will be the triangular area satisfying the inequalities

$$(11.2) \qquad x_i \geqq 0 \qquad i = 1, 2$$

and

$$(11.3) \qquad x_1 + x_2 \leqq 50.$$

If these vectors are (a^1, a^2), then $X = X^{12}$ will be the triangle satisfying (11.2) as well as

$$(11.4) \qquad 2x_1 + x_2 \leqq 100.$$

If these vectors are (a^2, a^1), then $X = X^{21}$ will be the triangle satisfying (11.2) as well as

$$(11.5) \qquad x_1 + 2x_2 \leqq 100.$$

Finally, if these vectors are (a^2, a^2), then $X = X^{22}$ will be the triangle satisfy-

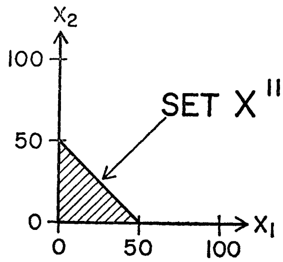

FIGURE 1

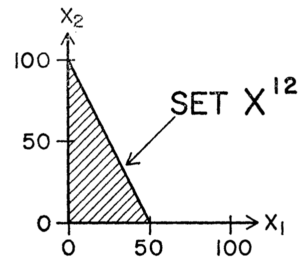

FIGURE 2

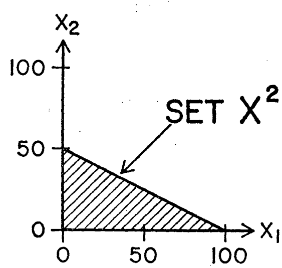

FIGURE 3

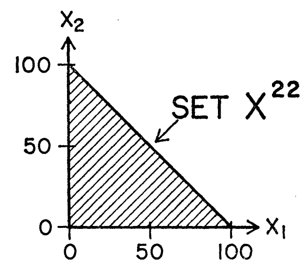

FIGURE 4

ing (11.2) as well as

(11.6) $$x_1 + x_2 \leqq 100$$

(see Figures 1 to 4).

Now suppose that the basic probability matrix $\{r_{km}\}$ of the game is as follows:

TABLE 22

	$c_2 = a^1$	$c_2 = a^2$
$c_1 = a^1$	$r_{11} = \frac{1}{2}\epsilon$	$r_{12} = \frac{1}{2} - \frac{1}{2}\epsilon$
$c_1 = a^2$	$r_{21} = \frac{1}{2} - \frac{1}{2}\epsilon$	$r_{22} = \frac{1}{2}\epsilon$

Here $r_{km} = \text{Prob}\ (c_1 = a^k \text{ and } c_2 = a^m)$ while ϵ is a very small positive number which for practical purposes can be taken to be zero.

Accordingly, both players i will use the following conditional probabilities (subjective probabilities):

TABLE 23

	$c_j = a^1$	$c_j = a^2$
$\text{Prob}(c_j = a^m \mid c_i = a^1)$	$p_{11} = q_{11} = \epsilon$	$p_{12} = q_{21} = 1 - \epsilon$
$\text{Prob}(c_j = a^m \mid c_i = a^2)$	$p_{21} = q_{12} = 1 - \epsilon$	$p_{22} = q_{22} = \epsilon$

That is, each player will assume a virtually complete negative correlation between his own attribute vector and the other player's: he will expect the other player to be free from the side-payment obligation if he himself is subject to it, and will expect the other player to be subject to this obligation if he himself is free from it.

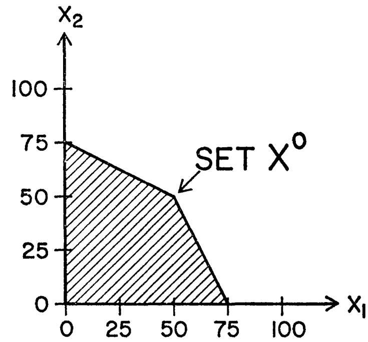

FIGURE 5

In constructing the normal form of this game, we shall take $\epsilon = 0$. Hence, the players will have $\frac{1}{2}$ probability of being able to select a payoff vector $x = x^*$ from the set X^{12}, and have $\frac{1}{2}$ probability of being able to select a payoff vector $x = x^{**}$ from the set X^{21}. Thus, their payoff vector from the normal-form game as a whole will be

$$(11.7) \qquad x = \tfrac{1}{2}x^* + \tfrac{1}{2}x^{**} \quad \text{with} \quad x^* \in X^{12}, \qquad x^{**} \in X^{21}.$$

The set $X^0 = \{x\}$ of all such vectors x is the area satisfying condition (11.2) as well as the two further inequalities:

$$(11.8) \qquad 2x_1 + x_2 \leqq 75$$

$$(11.9) \qquad x_1 + 2x_2 \leqq 75.$$

(See Figure 5.)
Thus the normal form $\mathfrak{N}(G)$ of this game will be a two-person bargaining game in which the players can choose any payoff vector $x = (x_1, x_2)$ from the set X^0, and obtain the payoffs $x_1 = 0$ and $x_2 = 0$ if they cannot agree on which particular vector x to choose.

Any solution concept satisfying the usual Efficiency and Symmetry Postulates will yield the solution $x_1 = x_2 = 50$. This, however, means that if case (a^1, a^2) arises, then the players will obtain the payoffs $x_1^* = 0, x_2^* = 100$, whereas if case (a^2, a^1) arises, then they will obtain the payoffs $x_1^{**} = 100$ and $x_2^{**} = 0$. This is so because, in view of condition (11.7), the only way we can obtain the vector $x = (50, 50)$ is by choosing $x^* = (0, 100)$ and $x^{**} = (100, 0)$.

If we interpret the normal form of the game in the usual way, and assume that the players will choose their strategies and will agree on the outcome *before* they know whether case (a^1, a^2) or case (a^2, a^1) will obtain, then as a solution of the normal form of the game this result makes very good sense. It means that in accordance with the Efficiency Postulate the players will try to *minimize* the expected value of the amount they would have to pay out to outsiders, and therefore will agree to assign a zero payoff to that player who would be under an obligation to hand over half of his payoff to the secret organization in question.

However, if we stick to the original definition of the game, under which the players can choose their strategies and can conclude agreements only *after* each player has observed the value of his own attribute vector c_i, then this result is highly counter-intuitive. This is so because in this game there is near-perfect negative correlation between the two players' attribute vectors; so that each player will not only know the value of his own attribute vector but will also be able to infer with virtual certainty what the value of the other player's attribute vector is. Hence, even though formally the game is a game with *incomplete* information, it will always essentially reduce to a two-person bargaining game with *complete* information.

More particularly, if case (a^1, a^2) arises, then this game G will be essentially equivalent to a Nash-type bargaining game G^{12} where the payoff space is the set $X = X^{12}$, and where this fact is known to both players. Likewise, if case (a^2, a^1) arises, then game G will be essentially equivalent to a Nash-type bargaining game G^{21} where the payoff space is the set $X = X^{21}$, and where again this fact is known to both players. Yet, it would be quite counter-intuitive to suggest that in game G^{12} player 1 would agree to a solution giving him a zero payoff $x_1^* = 0$, or that in game G^{21} player 2 would agree to a solution giving him a zero payoff $x_2^{**} = 0$.

In terms of the corresponding Bayesian game G^*, we can state our argument as follows. If the players can reach an agreement *before* a chance move decides whether case (a^1, a^2) or case (a^2, a^1) will occur, then each player may be quite willing to agree to an arrangement under which he would obtain \$100 in one of the two possible cases, and would obtain \$0 in the other case, both cases having the same probability $\frac{1}{2}$. But if the players can conclude an agreement only *after* the outcome of this chance event has come to be known to both of them, then neither player will voluntarily accept a payoff of \$0 in order that the other player can obtain all the \$100—even if such an arrangement would minimize the two players' total payment to the secret organization.

For definiteness let us assume that in actual fact in games G^{12} and G^{21} the players will use the Nash solution [4]. Then in G^{12} their payoffs will be $x_1^* = 25$ and $x_2^* = 50$, whereas in G^{21} their payoffs will be $x_1^{**} = 50$ and $x_2^{**} = 25$. Hence, by equation (11.7) the players' expected payoffs from game G^* as a whole will be $x_1 = \frac{1}{2}x_1^* + \frac{1}{2}x_1^{**} = 37.5$ and $x_2 = \frac{1}{2}x_2^* + \frac{1}{2}x_2^{**} = 37.5$.[2]

Of course, for game G^* as a whole, the vector $x = (37.5, 37.5)$ is an inefficient payoff vector because it is strongly dominated, e.g., by the payoff vector $x = (50, 50)$ previously considered. Yet in the case of game G^* the use of an inefficient solution is justified by the fact that G^* is not a cooperative game in the full sense of this term, precisely because the players are not allowed to conclude agreements at the very beginning of the game, before any chance move or personal move has occurred—but rather can conclude agreements only *after* the chance move deciding between cases (a^1, a^2) and (a^2, a^1) has been completed. Thus, by using the inefficient payoff vector $x = (37.5, 37.5)$ as the solution for game G^*,

[2] In terms of the terminology to be introduced in Part III, the bargaining game G under discussion can be *decomposed* into G^{12} and G^{21} as component games. Hence, according to Postulate 3 of Part III, whatever solution concept we may wish to use, the solution of game G must be equivalent to using the solution of game G^{12} or of game G^{21} depending on which of these two games represents the game really played by the two players.

we do not violate the principle that each fully cooperative game should have an efficient solution.[3, 4]

We can state our conclusion also as follows. Most games discussed in the game-theoretical literature are *games with immediate commitment*, in the sense that each player is free to commit himself to some particular strategy *before* any chance move or personal move has taken place. (In the special case where the game is a cooperative game, the players are also free to enter into fully binding and enforceable agreements at that point.)

In contrast, the Bayesian games G^* we are using in the analysis of games with incomplete information, must be interpreted as *games with delayed commitment*,[5] where the players can commit themselves to specific strategies (and can enter into binding agreements if the game they are playing is a cooperative game), only *after* certain moves have occurred in the game, viz., the chance moves determining the players' attribute vectors.[6]

By the very definition of the normal form of a game, von Neumann and Morgenstern's Normalization Principle [6] can apply only to games with immediate commitment. Hence, in general, we cannot expect that the normal form $\mathfrak{N}(G^*) = \mathfrak{N}(G)$ of a Bayesian game G^* will be a fully satisfactory representation of this game G^*—or of the incomplete-information game G Bayes-equivalent to G^*. This is the reason why in the analysis of these games we have found it necessary to replace von Neumann and Morgenstern's Normalization Principle by our weaker Semi-normalization Principle (Postulate 2 in section 7 of Part I of this paper).

REFERENCES

1. Gerard Debreu, "A Social Equilibrium Existence Theorem," *Proceedings of the National Academy of Sciences*, 38 (1952), 886–893.
2. John C. Harsanyi and Reinhard Selten, "A Generalized Nash Solution for Two-Person Bargaining Games with Incomplete Information," 1967 (unpublished).
3. John F. Nash, "Equilibrium Points in n-Person Games," *Proceedings of the National Academy of Sciences*, 36 (1950), 48–49.
4. ——, "The Bargaining Problem," *Econometrica*, 18 (1950), 155–162.
5. ——, "Non-Cooperative Games," *Annals of Mathematics*, 54 (1951), 286–295.
6. John von Neumann and Oskar Morgenstern, *Theory of Games and Economic Behavior*, Princeton: Princeton University Press, 1953.

[3] To be sure, game G itself if a fully cooperative game, yet it is a cooperative game with incomplete information. But our point has been that, as we replace this I-game G by a Bayes-equivalent C-game G^*, the fully cooperative nature of the game is lost since the players' attribute vectors c_i come to be reinterpreted as chance moves *preceding* their choices of strategies.

[4] Our present discussion has been restricted to a very special case of two-person bargaining games with incomplete information, viz., to the case where almost perfect correlation exists between the two players' attribute vectors c_1 and c_2. The problem of defining a suitable solution concept in the general case will be discussed in [2].

[5] Quite apart from this connection with Bayesian games, the general class of games with delayed commitment seems to have many other applications, and would deserve further study.

[6] Note that the corresponding Selten games G^{**} do not give rise to this problem: they can always be interpreted as games with immediate commitment and therefore can be adequately represented by their normal forms $\mathfrak{N}(G^{**})$.