

ϵ -Bayesian Solutions of Zero Sum Games in
Extensive Form

Jean-Pierre Ponsard

August 1974

WP-74-37

Working Papers are not intended for distribution outside of IIASA, and are solely for discussion and information purposes. The views expressed are those of the author, and do not necessarily reflect those of IIASA.

ϵ -BAYESIAN SOLUTIONS OF ZERO SUM GAMES
IN EXTENSIVE FORM

Jean-Pierre Ponsard

International Institute for Applied Systems Analysis
Laxenburg, Austria

Revised August 1974

ABSTRACT

This paper studies various extensions of the minimax principle to select unique optimal behavioral strategies in zero sum games in extensive form. Formal comparisons of the different extensions are carried out in the context of zero sum sequential games with incomplete information. There it appears that the various extensions are related to different approaches to the exploitation of the opponent's mistakes.

1. Introduction

Our basic motivation is to understand the rationale associated with an optimal behavioral strategy in a zero sum game defined in extensive form.

The traditional approach to solve such games is to reduce them to their normal form and derive the optimal behavioral strategies from the optimal mixed strategies. This procedure has the conceptual drawback of not providing much insight for the extensive solution. Furthermore it has the technical drawback of enlarging the dimension of the problem considerably (while the dimension of the behavioral strategy set grows linearly with the number of information sets in the game tree, the dimension of the mixed strategy set grows exponentially), as a result the optimal mixed strategies are in general non unique though the optimal behavioral strategy might be unique. (This of course generates some frustration after tedious computations!)

Some recent approaches, however, propose direct procedures to solve games in extensive form [P1, H-R, W]. All these approaches use the fact that optimal behavioral strategies are "Bayesian" in the sense that they maximize the conditional expectations at each information set given the other player's optimal strategy. Now, there is a difficulty in this fact, it concerns non optimal play. Indeed after a non optimal move, conditional expectations might not be mathematically defined. The purpose of this paper is to extend the Bayesian interpretation in this precise case.


Three ϵ -Bayesian extensions will be studied, first a "normal" one in which each pure strategy has to be played with a small probability ϵ and an "extensive" one in which each personal information set has to be reached with a small probability ϵ . These extensions were suggested by John Harsanyi. A third extension due to Melvin Dresher [D] will also be studied. Formal comparisons will be made in the framework of zero sum sequential games with incomplete information [P-Z].

2. The Main Ideas

This section presents the main ideas by means of examples. The first example is elementary; it demonstrates why non-uniqueness arises after a non optimal move. However, in this case, the non uniqueness is trivially solved. In the second example, each extension will generate a different optimal reply.

2.1 An Elementary Example

Consider the following zero sum game in extensive form in which Player 1 is the maximizer and Player 2 the minimizer.



The normal form of the game is

| | | |
|---|---|----|
| | c | d |
| a | 1 | -1 |
| b | 0 | 0 |

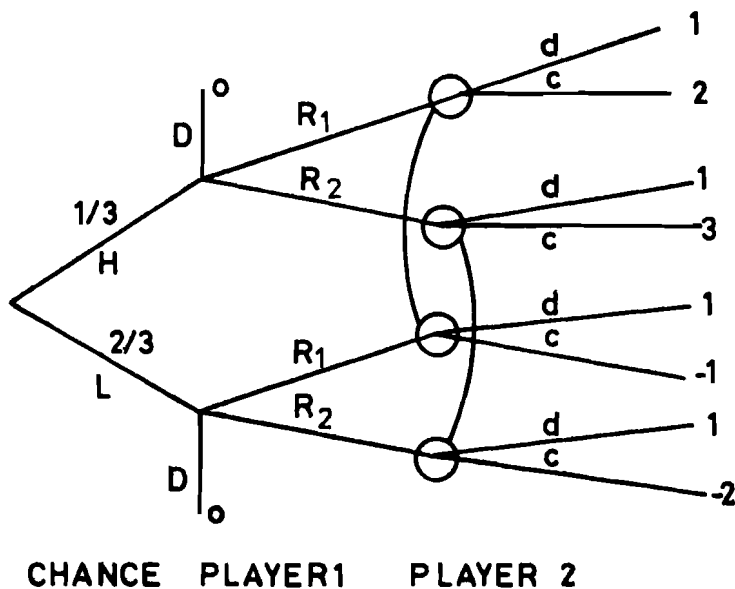
Player 2's optimal strategy is non unique and may be any convex combination of $(0c + 1d)$ ^{1/} and $(1/2c + 1/2d)$. Clearly if Player 2 wants to maximize his conditional expectation if he gets to play he should select $(0c + 1d)$ and any intuitive rationale would call this his optimal reply. The reason why the normal form fails to identify this unique solution is

^{1/} $(\alpha c + (1-\alpha)d)$ means that Player 2 selects move c with probability α and move d with probability $(1-\alpha)$. This notation is used consistently.

simply because it is concerned with unconditional expectations. Whatever Player 2 does after move a contributes nothing to the unconditional payoff so long as move a remains a mistake.

2.2 The Main Example

The zero sum game described by the following game tree may be interpreted as a one stage simplified poker. Player 1 receives one card which may be low (L) with probability $2/3$ or high (H) with probability $1/3$. Then he may drop (D), raise 1 unit (R1) or raise two units (R2). If Player 1 raised then Player 2 may drop (d) or call (c). There is one unit in the pot at the beginning of the game and the payoffs have been computed so that Player 1 is the maximizer.



The normal form of this game is already quite large:

| Player 2 Player 1 | dd ²⁾ | dc | cd | cc |
|----------------------|------------------|------|------|------|
| DD ¹⁾ | 0 | 0 | 0 | 0 |
| DR1 | 2/3 | 2/3 | -2/3 | -2/3 |
| DR2 | 2/3 | -4/3 | 2/3 | -4/3 |
| R1D | 1/3 | 1/3 | 2/3 | 2/3 |
| R1R1 | 1 | 1 | 0 | 0 |
| R1R2 | 1 | -1 | 4/3 | -2/3 |
| R2D | 1/3 | 1 | 1/3 | 1 |
| R2R1 | 1 | 5/3 | -1/3 | 1/3 |
| R2R2 | 1 | -1/3 | 1 | -1/3 |

Player 1's optimal mixed strategy is unique: (2/3 R2D, 1/3 R2R2).

Player 2's optimal mixed strategy set has four extremal points

MS1: (1/3dd, 0dc, 1/3cd, 1/3cc)

MS2: (0dd, 1/3dc, 2/3cd, 0cc)

MS3: (1/6dd, 1/3dc, 1/2cd, 0cc)

MS4: (1/2dd, 0dc, 1/6cd, 1/3cc)

However its optimal behavioral strategy set has only 2

BS1: (1/2d + 1/2c | R1) and (2/3d + 1/3c | R2)

BS2: (1/3d + 2/3c | R1) and (2/3d + 1/3c | R2)


(It is early seen that MS1 and MS2 collapse into BS1 and, MS3 and MS4 into BS2).

1) (D,D) stands for (drop with a high card, drop with a low card) and so on.

2) (d,d) stands for (drop if Player 1 raises 1, drop if Player 1 raises 2) and so on.

Hence this second example is similar to the first one in the sense that non uniqueness occurs after a non optimal move, namely move R1. BS1 and BS2 are the two extremal behavioral strategies (just as $(1/2c + 1/2d)$ in example 1) such that R1 remains a mistake. Note that Player 2's optimal behavioral strategies guarantee that he should not pay more than $5/3$ if the card is H and 0 if it is L. Now the exploitation of the mistake should somewhat improve these security levels. How to improve them is another story as we shall see.

Let us first get a clearer picture of the situation by drawing a diagram of Player 2's conditional expectation given R1.



PLAYER 2'S CONDITIONAL EXPECTATION
GIVEN R1

FIGURE 1

Mathematically speaking this expectation conditional on Player 1's optimal strategy does not exist since R1 should never be played. Taking a normal point of view any convex combination of BS1 and BS2 makes good sense since it guarantees that R1 is dominated and should not be played. It operates like a threat and the precise exploitation of the mistake becomes somewhat irrelevant. Now in the game in extensive form the rationale should be just the reverse: threats ought to be explicitly modelled as at what time they can be made binding and known to the opponent, whereas the precise exploitation of a mistake becomes interesting since it may very well be committed and known.

Coming back to our example, the exploitation of the mistake R1 is controversial since it would require the knowledge of the a posteriori probability $\text{Prob} (H | R1)$. Taking a Bayesian view, once an assumption about $\text{Prob} (H | R1)$ is made, Player 2's best reply is determined. However note that any convex combination of BS1 and BS2 can only be consistent with $\text{Prob} (H | R1) = 2/3$. This in turns implies a great deal about how R1 may be played, namely $\text{Prob} (R1 | H) = 4 \text{ Prob} (R1 | L)$. So that

$$\begin{aligned} \text{Prob} (H | R1) &= \frac{\text{Prob} (H) \cdot \text{Prob} (R1 | H)}{\text{Prob} (H) \cdot \text{Prob} (R1 | H) + \text{Prob} (L) \cdot \text{Prob} (R1 | L)} \\ &= \frac{1/3 \cdot 4 \text{ Prob} (R1 | L)}{(1/3 \cdot 4 + 2/3 \cdot 1) \text{ Prob} (R1 | L)} \\ &= 2/3. \end{aligned}$$

Now comes a difficult dilemma:

- (i) either Player 2 wants to select his behavioral strategy among the convex combinations of BS1 and BS2 and be consistent with game theory but he has to admit that he is acting as if the mistake was committed in some very specific way,
- (ii) or Player 2 makes a behavioral assumption about how the mistake was committed, derives a conditional probability which may be different from $2/3$ and maximizes his expectation at this point but he has to admit that if Player 1 were to know Player 2's strategy then his behavioral assumption would not make sense (playing R1 might not be a mistake any longer).

We may refer to (i) as a "tricky opponent" and to (ii) as a "stupid opponent". Clearly in a practical setting one approach might be more appropriate than the other. Nevertheless since obviously there is not much to say along the "stupid line" we propose to pursue a little longer the "tricky line". We shall examine three extensions of the game solution concept each of which will ordinarily determine a unique behavioral strategy for Player 2.

2.2.1. Extension 1

In this extension we shall simply assume that Player 2 wants to maximize the expected penalty given the a priori probability on H or L of $(1/3, 2/3)$ but constrains himself

to the convex combinations of BS1 and BS2. This procedure generates a unique optimal behavioral strategy which clearly is BS2 (see figure 1).

We shall demonstrate later on that this extension is associated with the following assumption on Player 1's behavior:

Player 2 assumes that Player 1 has to make a mistake (move R1) when he gets a low card and when he gets a high card the probability of the mistake being at least ϵ .

2.2.2. Extension 2 (Dresher)

This extension consists in restricting Player 1's pure strategy set to those which are not best replies to Player 2's optimal strategy set and in restricting Player 2 to remain in his own optimal strategy set [D]. In this case we obtain the matrix game.

| Player 2 | | |
|----------|-----|-----|
| Player 1 | BS1 | BS2 |
| DD | 0 | 0 |
| DR1 | 1/2 | 5/9 |
| R1D | 0 | 0 |
| R1R1 | 1/2 | 1/3 |
| R1R2 | 5/9 | 1/3 |
| R2D | 0 | 0 |
| R2R1 | 1/2 | 5/9 |

and Player 2's unique optimal strategy is $(4/5 \text{ BS1} + 1/5 \text{ BS2})$. Player 1's "optimal mistakes" may be any convex combination of $(4/5 \text{ DR1} + 1/5 \text{ R1R2})$ or $(1/5 \text{ R1R2} + 4/5 \text{ R2R1})$. Note however

that for any of these combinations $\text{Prob}(R1 | H) = 4 \text{ Prob}(R1 | L)$ and so $\text{Prob}(H | R1) = 2/3$.

We may now determine the penalty associated with Player 2's strategy (see figure 1):

- if the card is H, Player 2 can guarantee to lose no more than $5/3$, but if R1 is played then he will obtain $1/5 \cdot 5/3 + 4/5 \cdot 3/2 = 23/15$ thus he will be better off by $5/3 - 23/15 = 2/15$.
- if the card is L, Player 2 can guarantee to lose no more than 0, but if R1 is played then he will obtain $-1/5 \cdot 1/3 + 4/5 \cdot 0 = -1/15$ thus he will be better off by $0 + 1/15 = 1/15$. Since the priori probability of H is $1/3$ and of L is $2/3$ he will be better off equally well in a priori expectation basis whether the card is H or L.

We shall demonstrate later on that this extension is associated with the following assumption on Player 1's behavior:

Player 2 assumes that Player 1 has to make a mistake (move R1) with a small probability ϵ , Player 1 may decide to make the mistake with a low or a high card, but the overall probability of the mistake has to be at least ϵ . Then Player 2's optimal strategy will generate the same expected penalty (on an a priori basis) whether the card is L or H.

2.2.3. Extension 3

Taking a strictly Bayesian view, Player 2 should select his strategy by taking into consideration only the payoffs associated with move R1, the a posteriori probability on

H and L and his conditional security levels. Following this line, we shall propose the following behavioral assumption:

Player 2 will select his strategy so as to maximize the expected penalty at the a posteriori probability distribution and Player 1 will select this a posteriori probability distribution so as to minimize the penalty.

In the matrix form the game is

| | | |
|----------------------|-----|------|
| Player 1 Player 2 | d | c |
| H | 2/3 | -1/3 |
| L | -1 | 1 |

in which the entry (H, d) is computed as follows given that R1 is played the payoff with a high card is 1, since Player 2 can guarantee himself 5/3, the penalty is $5/3 - 1 = 2/3$.

The other entries are computed similarly.

Player 2's optimal behavioral strategy is $(4/9d + 5/9c)$ if R1 is played. Note that this coincides with $(2/3BS1 + 1/3BS2)$.

The best a posteriori probability for Player 1 is:

$\text{Prob}(H | R1) = 2/3$. Finally note that the conditional penalties incurred by Player 1 are identical and equal to 1/9.

Indeed, we shall demonstrate later on that this is one of the significant properties of this extension. We shall also prove that an equivalent behavioral assumption to the one underlined above is:

Player 2 assumes that Player 1 has to make a mistake with a small probability ϵ , Player 1 may decide to make the mistake with a low or high card but the overall probability of reaching the information set R_1 has to be at least ϵ .

3. ϵ -Bayesian Solutions for Zero Sum Sequential Games with Incomplete Information.

We shall now formalize the extensions presented in section 2 for the class of zero sum sequential games with incomplete information [P-Z].

3.1 Definition of the Game and its L-P Formulation

The game consists of four steps:

Step 0: chance chooses a move $k \in K$ according to a probability distribution $p^0 = (p_k^0)_{k \in K}$.

Player 1 is informed of the move chosen by chance, Player 2 is not.

Step 1: Player 1 chooses a move $i \in I$. Player 2 is informed of the move chosen by Player 1.

Step 2: Player 2 chooses a move $j \in J$.

Step 3: Player 1 receives an amount a_{kj}^i (a real number) from Player 2.

(We assume that K, I, J are all finite sets).

Let $x = (x_i^k)_{k \in K, i \in I}$ be a behavioral strategy for Player 1 such that $x_i^k = \text{Prob}(i | k)$ and $y = (y_j^i)_{i \in I, j \in J}$ be a

behavioral strategy for Player 2 such that $y_j^i = \text{Prob}(j | i)$.

Then Player 1 and Player 2 optimal behavioral strategies are the solutions the linear programs LP1 and LP2 respectively. [P2].

(See also [H-R] for a general formulation of Bayesian equilibrium in non-zero sum games with incomplete information).

| | |
|--|--|
| $\begin{aligned} \text{Max} \quad & \sum_{i \in I} u_i \\ \text{for } j \in J, i \in I \quad & u_i - \sum_{k \in K} a_{kj}^i p_k^o x_i^k \leq 0 \\ \text{for } k \in K \quad & \sum_{i \in I} x_i^k = 1 \\ \text{for } k \in K, i \in I \quad & x_i^k \geq 0 \end{aligned}$ <p style="text-align: center;">(LP1)</p> | $\begin{aligned} \text{Min} \quad & \sum_{k \in K} p_k^o v_k \\ \text{for } k \in K, i \in I \quad & v_k - \sum_{j \in J} a_{kj}^i y_j^i \geq 0 \\ \text{for } i \in I \quad & \sum_{j \in J} y_j^i = 1 \\ \text{for } i \in I, j \in J \quad & y_j^i \geq 0 \end{aligned}$ <p style="text-align: center;">(LP2)</p> |
|--|--|

The optimal values of the variables $(v_k)_{k \in K}$ may be interpreted as Player 2's conditional security levels for each respective $k \in K$.

3.2 ϵ -Bayesian Solutions

Suppose that one knows the optimal $(\bar{v}_k)_{k \in K}$ then LP2 is decomposable and for each $i \in I$ the optimal $y^i = (y_j^i)_{j \in J}$ is a solution of the problem (denoted by LP2i):

$$\begin{aligned} k \in K \quad \bar{v}_k - \sum_{j \in J} a_{kj}^i y_j^i & \geq 0 \\ \sum_{j \in J} y_j^i & = 1 \\ j \in J \quad y_j^i & \geq 0 \end{aligned}$$

Now, the two linear programs LP1 and LP2 are dual of each other and the dual variables associated with the inequality

$$v_k - \sum_{j \in J} a_{kj}^i y_j^i \geq 0$$

is $p_k^o x_i^k$.

Thus if for some $i \in I$ and all $k \in K$, $\bar{x}_k^i = 0$ (that is: move i is non optimal for Player 1) then the constraints in the associated LP2i problem are not binding at the optimum and ordinarily this will result in non uniqueness for the optimal \bar{y}^i . The whole issue of resolving this non-uniqueness amounts to defining a proper objective function for LP2i. We shall now exhibit the objective functions associated with the three extensions presented in section 2.

3.2.1. Extension 1

The idea in extension 1 is that every pure strategy should be played with a small probability ϵ . In this context this amounts to impose the additional constraints in LP1:

$$i \in I, k \in K \quad x_i^k \geq \epsilon$$

Let $(\gamma_i^k)_{k \in K, i \in I}$ be the dual variables associated with these constraints. Consequently LP2i becomes (assuming ϵ small enough so that the optimal $(\bar{v}_p)_{k \in K}$ will remain unchanged)

$$\text{Min } \epsilon \left(\sum_{k \in K} - p_k^0 \gamma_i^k \right)$$

$$k \in K \quad \bar{v}_k - \gamma_i^k - \sum_{j \in J} a_{kj}^i y_j^i = 0$$

$$\sum_{j \in J} y_j^i = 1$$

$$j \in J \quad y_j^i \geq 0$$

$$k \in K \quad \gamma_i^k \geq 0$$

Hence γ_i^k may be interpreted as the conditional penalty of playing move i in state k and Player 2's optimal behavioral strategy will maximize the expected penalty given the a priori probability distribution on k ($p^0 = (p_k^0)_{k \in K}$) under the constraints that each conditional penalty remains positive.

3.3.2. Extension 2

The idea in extension 2 is that the overall probability of playing each move be a small probability ϵ . In this context this amounts to impose the additional constraints in LP1

$$i \in I \quad \sum_{k \in K} x_i^k \geq \epsilon$$

Letting $(\gamma_i)_{i \in I}$ be the dual variables associated with these constraints we obtain for LP2i:

$$\begin{aligned} & \text{Min} - \epsilon \gamma_i \\ & k \in K \quad p_k^0 (\bar{v}_k - \sum_{j \in J} a_{kj}^i y_j^i) - \gamma_i \geq 0 \end{aligned}$$

$$\sum_{j \in J} y_j^i = 1$$

$$j \in J \quad y_j^i \geq 0$$

Hence Player 2's optimal behavioral strategy may be seen as the solution of the matrix game $||b_{kj}||_{k \in K, j \in J}$ with $b_{kj} = p_k^0 (\bar{v}_k - a_{kj}^i)$. It is easily seen that this corresponds to Drescher's approach in which Player 1 may only randomize among the pure strategies which include to play move i in some

state k whereas Player 2 is constrained to his optimal strategy set. Note that for all states k in which a mistake is made,

$$p_k^o (\bar{v}_k - \sum_{j \in J} a_{kj}^i \bar{y}_j^i) = \bar{\gamma}_i .$$

3.2.3. Extension 3

The idea in extension 3 is that each personal information set in the game tree should be reached with at least a small probability ϵ . Thus the additional constraints in LP1 are

$$i \in I \quad \sum_{k \in K} p_k^o x_i^k \geq \epsilon$$

Letting $(\gamma_i)_{i \in I}$ be the dual variables associated with these constraints, we obtain for LP2i:

$$\begin{aligned} & \text{Min} - \epsilon \gamma_i \\ & k \in K \quad \bar{v}_k - \sum_{j \in J} a_{kj}^i y_j^i - \gamma_i \geq 0 \\ & \sum_{j \in J} y_j^i = 1 \\ & j \in J \quad y_j^i \geq 0 \end{aligned}$$

Then Player 2's optimal behavioral solution is the solution of the matrix game $\| \| c_{kj} \| \|_{k \in K, j \in J}$ in which $c_{kj} = \bar{v}_k - a_{kj}^i$.

3.3. Discussion

Given the preceding formalization it is quite easy to determine what might be called ϵ -Bayesian^{1/} solutions in simple examples. Indeed, as long as the number of states

^{1/} The term "Bayesian" is used to emphasize that these extensions rely on mutually consistent optimal strategies for the two players.

remains equal to two, a graphical analysis is ordinarily quite sufficient. Here is such a graphical analysis which should convey some intuitive feelings about the differences between the three extensions.




FIGURE 2

In figure 2, we see that extension 1 is consistent with any a posteriori probability within the line segment $[p_2, p_1]$ whereas extension 2 is only consistent with p_2 and extension 3 with p_3 . (If we measure the "trickiness" of Player 1 by the difference between the a priori and the a posteriori probability distribution this would suggest different levels of trickiness in Player 1's behavior associated with the different extension).

Note also that as p_0 varies inside the open line segment $(0, 1)$ (assume for simplicity of the argument that the $(\bar{v}_k)_{k \in K}$ remain constant) then the behavioral strategy associated with extension 1 varies in a discrete fashion, whereas it varies smoothly in extension 2 and remains constant in extension 3. Were p_0 to be an extreme point such as 0 then \bar{v}_{k_1} , becomes irrelevant and all three extensions coincide with BSO, the intuitive solution of example 1 in section 2.

Finally note that only extension 3 will guarantee that whatever the mistake it be penalized. That is, suppose p_0 is within $(0, p_1)$ then extensions 1 and 2 will not penalize the mistake in state k_1 whereas extension 3 will. Formally speaking, it is known that there exists an optimal strategy for Player 2 such that for every pure strategy of Player 1 which is not part of his optimal strategy, the associated expected payoff is strictly less than the value of the game $[A-B-B]$. Then extension 3 selects one such strategy.

So far we have been directly concerned with the case in which only one mistake is available to Player 1. If several mistakes were available then as long as the $(\bar{v}_k)_{k \in K}$ are unique the problem is decomposable so that mistakes do not interact with each other and the preceding analysis apply to each mistake. If the $(\bar{v}_k)_{k \in K}$ are non unique (and they are non unique on a subset Q of the simplex

$$P = \{p = (p_k)_{k \in K} \mid p_k \geq 0, \sum_{k \in K} p_k = 1\}$$

in which the dimension of Q is the dimension of P minus one) then LP2 is not decomposable and some problems may occur but we shall refrain from going into this mathematical singularity.

4. More Examples

We shall conclude this paper by one example designed to show the relevance of the ideas in more general extensive games.

A two stage zero sum sequential game with incomplete information.

There are two cards, say one white and one black, which will be presented in sequence to Player 1. The sequence is chosen at random. Player 1 will announce the color of the card and he may say the truth or lie. After each card is announced Player 2 will say correct or not correct. Each stage payoff is given by the table

| Player 2 Player 1 | correct | not correct |
|---------------------------------|---------|-------------|
| truth | 0 | 1 |
| lie | 1 | -1 |

Player 1 (Player 2) wants to maximize (minimize) the expected payoff over the two stages. It is assumed that intermediary payoffs are not revealed.

If the first stage only were to be played then the value would be $1/3$ and the Players' optimal behavioral strategies would be:

Player 1 would say the truth with probability $2/3$ and lie with probability $1/3$,

Player 2 would say correct with probability $2/3$ and incorrect with probability $1/3$.

Now in the two stage game the value is $2/3$ [P-Z]. Player 1's optimal behavioral strategy at the first stage is unchanged whereas at the second stage he should announce the color he did not announce at the first stage (in other words, Player 1 should remain consistent with himself: lie twice or say the truth twice). Player 2's optimal behavioral strategy on the other hand is non unique. He may merely repeat twice his one stage optimal behavioral strategy but this will not penalize mistakes by Player 1 such as announced sequences "white, white" or "black, black" which obviously cannot be true.

Applications of the extensions proposed in section 3 will lead to the following ϵ -Bayesian solutions:

1st stage: say correct with probability $1/3$ and incorrect with probability $2/3$.

2nd stage: if Player 1 reverses his announcements (after announcing "white" he says "black" or vice versa) then always say correct.

If Player 1 does not reverse his announcement and thus commits a mistake then according to extension 1 he should always say incorrect and according to extensions 2 or 3 he should say correct with probability $1/3$ and incorrect with probability $2/3$.

The interesting feature of this example is that to be in a position to exploit a possible mistake, Player 2 has to anticipate it. Were he to simply play his one stage minimax strategy at the first stage ($1/3$ correct + $2/3$ incorrect) then the penalization of the mistake would be impossible.

A similar example but more complicated so that it cannot be reproduced here is Kuhn's Simplified Poker [K]. This example may serve as a good conclusion: starting with a game in extensive form Kuhn's derives 12 extremal optimal mixed strategies for Player 1, this reduces to a one parameter family in terms of optimal behavioral strategies. Again this non uniqueness is associated with the possibility of non optimal play by Player 2 and may be resolved along the lines developed in this paper (the interested reader will verify that the ϵ -Bayesian solutions require the use of underbidding, namely passing with a high card. It is not hard to give an intuitive justification of such a behavior: it may attract the mistaken player into betting with a medium hand whereas his optimal strategy calls for passing).

REFERENCES

- [A-B-B] Arrow, K.J., E.W. Barankin and D. Blackwell,
"Admissible Points of Convex Sets". Contributions
to the Theory of Games, Vol. 2.
- [A-M] Aumann, R.J. and M. Maschler, "Some Thoughts on the
Minimax Principle", Management Science, Vol.18 No.5.
- [D] Dresher, M., Games of Strategy Theory and Applications,
Prentice Hall 1961.
- [H] Harsanyi, J., "The Tracing Procedure", University of
Bielefeld/IMW, No.15, May 1974.
- [H-R] Howson, J.T. and R.W. Rosenthal, "Bayesian Equilibria
of Finite Two-Person Games with Incomplete
Information", forthcoming in Management Science.
- [K] Kuhn, H.W., A Simplified Two Person Poker,
Contributions to The Theory of Games, Vol.1.
- [P-1] Ponssard, J.P., "Zero Sum Games with 'Almost'
Perfect Information", forthcoming in Management
Science.
- [P2] Ponssard, J.P., "A Note on the L-P Formulation of
Zero-Sum Sequential Games with Incomplete
Information", IIASA, Laxenburg, May 1974.
- [P-Z] Ponssard, J.P. and S. Zamir, "Zero Sum Sequential
Games with Incomplete Information", Journal of
Game Theory 2, No.2, 1973.
- [W] Wilson, R., "Computing Equilibria of Two Person
Games from the Extensive Form".