Recent Developments in Operational Research
Manju Lata Agarwal and Kanwar Sen (Eds)

4. Role of Information in the Interaction of Two Statisticians: Some game theoretic considerations

Bruno Bassan
Dipartimento di Matematica, Università "La Sapienza", Piozzale Aldo Moro 5, I-00185 Roma, Italy

Marco Scarsini
Dipartimento di Scienze, Università D'Annunzio, Viale Pindaro 42, I-65127 Pescara, Italy

Shmuel Zamir
Department of Statistics, Hebrew University of Jerusalem, Mount Scopus, Jerusalem 91905, Israel

Abstract: We consider situations in which two statisticians are faced with a decision, and the loss function of each of them depends also on the decision made by the other one, namely, we consider situations in which two statisticians are involved in a game.

We are interested in models that may fit 'real life' situations, and in which both statisticians may prefer to reject free information. This is in sharp contrast with the usual "rule" in statistics that asserts that an additional observation in a sample cannot be harmful.

After some general considerations on the problem, examples concerning Bayesian hypothesis testing and prediction will be shown.

1. Introduction

The relation between mathematical statistics and game theory has been deeply analyzed in the literature for about 50 years. Typically, a statistical decision problem can be seen as a game played by the Statistician against Nature. For example, a statistical decision problem with two actions and two states of nature can be represented as a game in which, say, the Statistician chooses the row and Nature chooses the column. Such a framework leads naturally to the notion of minimax estimators. For a review of the relationship between statistics and game theory, see Schwarz (1994) and references therein.

This approach was popular in the 1950s, but then it was neglected, mainly due to the conceptual difficulty in seeing Nature as an agent playing strategically against the Statistician. It is not even clear what the payoff for Nature in such a game should be. However, minimax estimators have drawn renewed attention in the recent years. See for instance Donoho and Johnstone (1994), (1996), (1998).

It is significant that the only game theoretical concept pertaining in Statistics is that of minimax. This concept which captures the idea of 'security level' or 'the worst scenario', depends neither on Nature's payoffs nor on Nature's 'strategic behavior'. It simply assumes that Nature acts so as to increase the Statistician's losses.

In Bayesian decision problems Nature does not play strategically, but it chooses among possible actions with a probability which is in general not fully known to the Statistician. For example, the two-actions-two-states case can be modeled as follows. Call s_1, s_2 the two states, and let $(\Omega, \mathcal{F})$ be a measurable space, on which the following random variables are defined: $S: \Omega \to \{s_1, s_2\}$ amd $\Theta: \Omega \to [0, 1]$. We may think of $\omega \in \Omega$ as a pure strategy for nature. The choice of $\omega \in \Omega$ will induce a choice of a state $s_i = S(\omega)$ and of $\theta = \Theta(\omega)$. The value θ will be interpreted as the probability with which s_1 is chosen. Therefore a mixed strategy for Nature is a probability $\mathbb{P}$ on $(\Omega, \mathcal{F})$ which satisfies the following constraints:

$$\mathbb{P}\left(\{\omega : S(\omega) = s_1\}/\Theta(\omega) = \theta\right) = \theta \quad a.\ s., \quad \theta \in [0, 1]$$

Let π denote the law of Θ. It follows that

$$\mathbb{P}(S = s_1) = \int_\Omega \Theta(\omega)\, \mathbb{P}(d\omega) = \int_0^1 \theta\, \pi(d\theta).$$

A sequence of random observations $X_1, X_2, \ldots$ will be defined on $(\Omega, \mathcal{F}, \mathbb{P})$. This will allow to define the conditional probabilities $\mathbb{P}(.| X_1, \ldots, X_n)$, $n \in \mathbb{N}$, which in turn will induce what in game theoretical terminology may be called the updated beliefs of the Statistician about the values $\Theta(\omega)$ and $S(\omega)$ chosen by Nature. In problems of hypothesis testing or estimation, what is relevant for the Statistician is Θ, whereas in binary prediction problems what is relevant is S.

In this note we deal with decision problems that involve two interacting statisticians. These models can be conveniently described by a game where Nature chooses $\omega \in \Omega$ (thus selecting a state), and each of the two statisticians chooses an action. The payoff (= – loss) of a statistician depends not only on the state of nature, but also on the action of the other statistician.

Some real life situations may be described by the above model. For example, think of a commercial transaction consisting in the sale of a large quantity of items. The sale is going to be carried out only if the statisticians of both the seller and the buyer approve the quality of the items, relying on possibly separate samples. In this case, which reduces to a hypothesis testing problem for each of the statisticians, there are other relevant issues besides making the 'correct' decision. For example, if the items are not of good quality, but they are declared acceptable by the buyer, the seller might have a higher payoff by declaring the items acceptable as well, even if a second-type error is being made. Other situations where two or more statisticians are involved are the so-called inspection games. See Avenhaus, von Stengel and Zamir (1995).

In these interactive statistical decisions problems, the usual criteria for selecting among decision rules must be replaced by interactive criteria, such as playing a Nash equilibrium. (A Nash equilibrium is a strategy profile such that no player can profit from unilaterally deviating from his strategy in the profile.)

Thus, well-established concepts such as sufficient statistic loose part of their relevance, and 'paradoxical' phenomena may happen. For example, it might be possible that playing the unique Nash equilibrium, or the Nash equilibrium most advantageous to both players, involves making a decision which is not based on a sufficient statistic. Of course, there is nothing paradoxical in this, since interactive criteria must replace the usual ones. We use the word 'paradoxical' in order to stress the difference with the usual rules.

In this note, we present two examples of what we may call 'information paradox': It is possible that both statisticians ignore (or refuse to acquire, if given the choice) an additional, free observation.

It is well known that information refusal cannot happen in usual statistical decision problems. In fact, consider a Bayesian Parametric problem, and assume for simplicity that only two actions are available, say a_0 and a_1. Let Θ be the (random) parameter about which inference is to be made and let $W(a, \theta)$ denote the loss function. Let $X_1, X_2, \ldots$ be observations whose law depends on Θ, and let

$$\rho(a, \mathbf{x}_n) = \mathbb{E}(W(a, \Theta) \,|\mathbf{X}_n = \mathbf{x}_n)$$

be the (posterior) expected loss after an n-sample x_n is observed. Clearly, for $j \in \{\mathbf{0.1}\}$, we have that $n \to \rho(a_j, \mathbf{X}_n)$ is a martingale. For each $\mathbf{x}_n$, consider the Bayes action $a^*_{\mathbf{x}_n}$ defined implicitly by

$$\rho(a^*_{\mathbf{x}_n}, \mathbf{x}_n) = \min \{\rho(a_0, \mathbf{x}_n), \rho(a_1, \mathbf{x}_n)\},$$

and consider Bayes rule δ^*_n given by $\delta^*_n(\mathbf{x}_n) = a^*_{\mathbf{x}_n}$. Clearly, $n \to \rho(\delta^*_n(\mathbf{X}_n), \mathbf{X}_n)$ is a supermartingale, being the minimum of two martingales. It follows that the Bayes risk associated to the Bayes rule δ^*_n

$$r_n(\delta^*_n) = \mathbb{E}(W(\delta^*_n(\mathbf{X}_n), \Theta)) = \mathbb{E}(\rho(\delta^*_n(\mathbf{X}_n), \mathbf{X}_n))$$

is decreasing in n. Hence the Statistician wants as many observations as possible Similar arguments apply also in Bayesian prediction problems (see Section 3 below).

The previous argument shows that the positive role of information in usual Bayesian statistical decision problems is a consequence of the fact that the Statistician chooses the action that minimizes the expected posterior loss. In interactive decision problems, this is no longer true, since the Statistician is concerned also by the decisions of the other statisticians and may thus choose for instance an action which is compatible with a Nash equilibrium. In such situations, there is no guarantee that the Statistician desires as much information as possible. Thus, the two examples below, a Bayesian hypothesis testing problem and a Bayesian prediction one, in which we see this phenomenon of information refusal, should not be considered surprising, although they are counterintuitive at first glance.

Games in which Nature chooses a state with a random probability are considered in Bassan, Scarsini and Zamir (1998), where the relation between uniqueness of

Pareto optima and positive role of information is investigated. Other references related to the information paradox can be found in the above paper.

2. Bayesian Testing

Let Θ be a (discrete) random variable taking values in the space $\{\theta_0, \theta_1\}$. We want to test the simple hypothesis

$$H_0\colon \Theta = \theta_0 \quad \text{versus} \quad H_1 : \Theta = \theta_1.$$

Consider the following bimatrices of payoffs:

$$\begin{array}{c|c|c|} & a_1 & a_0 \\ \hline a_1 & -1,-1 & -1,-1 \\ \hline a_0 & -1,-1 & 0,0 \\ \hline \end{array} \quad \text{if } \Theta = \theta_0; \quad \begin{array}{c|c|c|} & a_1 & a_0 \\ \hline a_1 & 0,0 & -1,-1 \\ \hline a_0 & -1,-1 & -1,-1 \\ \hline \end{array} \quad \text{if } \Theta = \theta_1. \tag{1}$$

The above matrices might arise in the following context: a seller and a buyer send each a statistician to inspect some items being sold. Unless both statisticians make the right decision each of the parties incurs in a unit loss (since either a bad item is being sold or the transaction is canceled.)

For convenience, we add 1 to every entry in (1). Thus, we work with

$$\begin{array}{c|c|c|} & a_1 & a_0 \\ \hline a_1 & 0,0 & 0,0 \\ \hline a_0 & 0,0 & 1,1 \\ \hline \end{array} \quad \text{if } \Theta = \theta_0; \quad \begin{array}{c|c|c|} & a_1 & a_0 \\ \hline a_1 & 1,1 & 0,0 \\ \hline a_0 & 0,0 & 0,0 \\ \hline \end{array} \quad \text{if } \Theta = \theta_1. \tag{2}$$

From now on, assume that $\theta_0 = \frac{1}{4}$ and $\theta_1 = \frac{3}{4}$. Assume also that the prior law of Θ, agreed upon by both players, is specified by

$$\mathbb{P}\left(\Theta = \frac{1}{4}\right) = \frac{3}{16} = \pi_0; \quad \mathbb{P}\left(\Theta = \frac{3}{4}\right) = \frac{13}{16} = 1 - \pi_0.$$

Assume that, given $\Theta = \theta$, the random variables $X_{\text{I}}, X_{\text{II}}, Y$ are i.i.d. Bernoulli with parameter θ. The game is as follows:

1. Each statistician first takes privately a sample of size 1 (X_{I} and X_{II}, respectively).
2. A sample of size 1, namely Y, is observed by both players.
3. Each statistician chooses his action.

In the first step, although the information received by each of the statisticians is private, it is common knowledge that it has been acquired. Namely it is common knowledge that Statistician I observed the value of X_{I} and Statistician II observed the value of X_{II}. Thus, using the jargon common in game theory (after Harsanyi (1967/1968)), we will say that player I can be of two types: I_0, if $X_{\text{I}} = 0$, and I_1, if $X_{\text{I}} = 1$. Similarly for player II. We write $\mathbb{P}_{\text{I}_0}(.)$ for $\mathbb{P}(.|X_{\text{I}} = 0)$, etc.

The above procedure involving the types might be seen as a way to model the fact that the statisticians possess different, private information. Notice that private information plays a crucial role in this example. In fact, it is proved in Bassan, Scarsini and Zamir (1998) that in games with a payoff structure like the one in (2), if the players have the same information they want as much information as possible about Θ.

We will see that the statisticians may or may not want that further information be released, depending on their type.

A strategy profile in this game is a string of 8 actions: the first two are the actions taken by I_0 if $Y = 0$ and $Y = 1$, respectively, and so on. Several equilibria emerge; among these:

1. Each player chooses a_1 if he thinks that θ_1 is more likely than θ_0, given the information available, including Y; this yields

$$\overset{(I_0)}{a_0 a_1}\ \overset{(I_1)}{a_1 a_1}\ \overset{(II_0)}{a_0 a_1}\ \overset{(II_1)}{a_1 a_1} \tag{3}$$

2. The same as above, but not taking into account Y; this yields

$$\overset{(I_0)}{a_1 a_1}\ \overset{(I_1)}{a_1 a_1}\ \overset{(II_0)}{a_1 a_1}\ \overset{(II_1)}{a_1 a_1} \tag{4}$$

In the following subsection we will show that:

(A) Each of the above strategy profiles is a Nash equilibrium;

(B) The payoffs for I_0 are $\frac{237}{352}$ if (3) is played, and $\frac{13}{22} < \frac{237}{352}$ if (4) is played. Hence, if given the choice, player I_0 would prefer that the additional information Y be revealed;

(C) The payoffs for I_1 are $\frac{195}{224}$ if (3) is played, and $\frac{39}{42} > \frac{195}{224}$ if (4) is played. Hence, if given the choice, player I_1 would prefer that the additional information Y not be revealed.

2.1 Computations

First observe that

$$\mathbb{P}_{\mathrm{I}_0}(\Theta = \theta_1) = \frac{\mathbb{P}(\Theta = \theta_1)\,\mathbb{P}(X_{\mathrm{I}} = 0/\Theta = \theta_1)}{\mathbb{P}(X_{\mathrm{I}} = 0)}$$

$$= \frac{(1-\pi_0)(1-\theta_1)}{(1-\pi_0)(1-\theta_1) + \pi_0(1-\theta_0)} = \frac{13}{22}.$$

Similarly, one finds $\mathbb{P}_{\mathrm{I}_1}(\Theta = \theta_1) = \frac{39}{42}$. Furthermore.

$$\mathbb{P}_{\mathrm{I}_1}(\Theta = \theta_1/Y = 0) = \mathbb{P}(\Theta = \theta_1/X_{\mathrm{I}} = 0, Y = 0)$$

$$= \frac{(1-\pi_0)(1-\theta_1)^2}{(1-\pi_0)(1-\theta_1)^2 + \pi_0(1-\theta_0)^2} = \frac{13}{40}$$

Similarly, one finds

$$\mathbb{P}_{\mathrm{I}_1}(\Theta = \theta_1/Y = 1) = \mathbb{P}_{\mathrm{I}_1}(\Theta = \theta_1/Y = 0) = \frac{13}{16}$$

and $$\mathbb{P}_{\mathrm{I}_1}(\Theta = \theta_1 / Y = 1) = \tfrac{117}{120}$$

Let us compute now the expected payoff of I_0 if (3) is played:

$$\begin{aligned}
\mathbb{P}_{\mathrm{I}_0}(\Theta = \theta_0, X_{\mathrm{II}} = 0, Y = 0) &= \mathbb{P}_{\mathrm{I}_0}(\Theta = \theta_1, Y = 1) \\
&= \mathbb{P}_{\mathrm{I}_0}(\Theta = \theta_0)\, \mathbb{P}_{\mathrm{I}_0}(X_{\mathrm{II}} = 0, Y = 0/\Theta = \theta_0) \\
&\quad + \mathbb{P}_{\mathrm{I}_0}(\Theta = \theta_1)\, \mathbb{P}_{\mathrm{I}_0}(Y = 1/\Theta = \theta_1) \\
&= \frac{9}{22}(1 - \theta_0)^2 + \frac{13}{22}\theta_1 = \frac{237}{352}
\end{aligned}$$

If I_0 unilaterally deviates from (3) and plays $a_1 a_1$ (other moves are clearly not advantageous), then his payoff becomes

$$\begin{aligned}
\mathbb{P}_{\mathrm{I}_0}(\Theta = \theta_1, Y = 0, X_{\mathrm{II}} = 1) &+ \mathbb{P}_{\mathrm{I}_0}(\Theta = \theta_1, Y = 1) \\
&= \frac{13}{22}[\theta_1(1 - \theta_1) + \theta_1] = \frac{195}{352}
\end{aligned}$$

Hence, I_0 has no interest in deviating from (3).

Now, we repeat the same computations for player I_1. If (3) is played, then his expected payoff is

$$\begin{aligned}
\mathbb{P}_{\mathrm{I}_1}(\Theta = \theta_1, X_{\mathrm{II}} = 1, Y = 0) &+ \mathbb{P}_{\mathrm{I}_1}(\Theta = \theta_1, Y = 1) \\
&= \frac{39}{42}[\theta_1(1 - \theta_1) + \theta_1] = \frac{195}{224}
\end{aligned}$$

If I_1 deviates and plays $a_0 a_1$, then his payoff becomes

$$\begin{aligned}
\mathbb{P}_{\mathrm{I}_1}(\Theta = \theta_0, X_{\mathrm{II}} = 0, Y = 0) &+ \mathbb{P}_{\mathrm{I}_1}(\Theta = \theta_1, Y = 1) \\
&= \frac{3}{42}(1 - \theta_0)^2 + \frac{39}{42}\theta_1 = \frac{165}{224}.
\end{aligned}$$

Thus, (3) is a Nash equilibrium. It remains only to compute the payoffs in the equilibrium (4). The payoff for I_0 is

$$\mathbb{P}_{\mathrm{I}_0}(\Theta = \theta_1) = \frac{13}{22} < \frac{237}{352},$$

whereas the payoff of I_1 is

$$\mathbb{P}_{\mathrm{I}_1}(\Theta = \theta_1) = \frac{39}{42} > \frac{195}{224}.$$

3. Bayesian Prediction

In this section we consider an example in which two statisticians have to predict the outcome of the next observation. After some preliminaries, we review why a single statistician faced with a Bayesian prediction problem always wants as many observations as possible. Then we consider the case involving two statisticians. The resulting game has some similarities with the one concerning Bayesian hypothesis testing, and in particular the matrices of payoffs are the same. Here, though, we don't confine ourselves to the dichotomy zero-observations/one-observation, but we consider also partial signaling. This means that a random variable is revealed to the players, and this random variable may have an arbitrary correlation coefficient, ranging from 0 to 1, with a random variable exchangeable with the one to be predicted. Thus a family of games emerges, indexed by the correlation coefficient. For each of these games, we will consider several equilibria, and we will see that, among these equilibria, the most advantageous for the players involves ignoring the signal, no matter what its intensity is.

Let Θ be a random variable distributed according to a Beta (α, β), and assume that $X_1, X_2, X_3 \ldots$ is a sequence of Bernoulli random variables, i.i.d. conditionally on Θ, with $\mathbb{P}(X_{n+1} = 1/\Theta = \theta) = \theta$.

3.1 One Statistician

A statistician has to predict X_{n+1}, after observing $X_1, \ldots, X_n$. The space of actions is formed by two points only, $\mathscr{A} = \{a_0, a_1\}$. The utility (= – loss) function is

$$u(a_j, X_{n+1}) = \begin{cases} 1 \text{ if } X_{n+1} = j, \\ 0 \text{ if } X_{n+1} \neq j. \end{cases}$$

Let us compute expected utilities:

$$\begin{aligned} \mathbb{E}(u(a_1, X_{n+1}) | X_1, \ldots, X_n) &= \mathbb{P}(X_{n+1} = 1 | X_1, \ldots, X_n) \\ &= \mathbb{E}\,(\mathbb{P}(X_{n+1} = 1 | X_1, \ldots, X_n, \Theta)\,| X_1, \ldots, X_n) \\ &= \mathbb{E}(\Theta \mid X_1, \ldots, X_n). \end{aligned}$$

Similarly, one shows that

$$\mathbb{E}\,(u(a_0, X_{n+1}) | X_1, \ldots, X_n) = \mathbb{E}\,(1 - \Theta \mid X_1, \ldots, X_n).$$

Hence, a rational decision maker, by taking n observations, achieves

$$M_n := \max\,\{\mathbb{E}\,(\Theta | X_1, \ldots, X_n), \mathbb{E}\,(1 - \Theta \mid X_1, \ldots, X_n)\}.$$

Since $n \to \mathbb{E}(\Theta | X_1, \ldots, X_n)$ is a martingale, we see that $n \to M_n$ is a submartingale. Hence $n \to \mathbb{E}(M_n)$ is increasing. It follows that a single decision maker, faced with this Bayesian prediction problem, wants as many observations as possible.

3.2 Two Statisticians

Assume that two statisticians must predict X_1. Assume that the reward to each statistician is 1 if both statisticians predict correctly the outcome of X_1, and 0 otherwise. Thus, the payoff matrices are as in (2), namely:

$$\begin{array}{c|c|c|} & a_1 & a_0 \\ \hline a_1 & 0,0 & 0,0 \\ \hline a_0 & 0,0 & 1,1 \\ \hline \end{array} \quad \text{if } X_1 = 0; \qquad \begin{array}{c|c|c|} & a_1 & a_0 \\ \hline a_1 & 1,1 & 0,0 \\ \hline a_0 & 0,0 & 0,0 \\ \hline \end{array} \quad \text{if } X_1 = 0. \qquad (5)$$

We assume that each of the statisticians has a private information: It is common knowledge that Statistician I observed the value of X_{I} and Statistician II observed the value of X_{II}, where X_{I}, X_{II}, X_1 are exchangeable random variables. Thus, Statistician I can be of two types: I_0, if $X_{\mathrm{I}} = 0$, and I_1, if $X_{\mathrm{I}} = 1$. Similar for Statistician II.

Assume that, in addition to the private observations, a public signal ξ is revealed to the players before they formulate their prediction. In order to describe this signal, consider first three independent Bernoulli random variables Y, W, Z such that

- Y, X_{I}, X_{II}, X_1 are exchangeable;
- W is independent of X_{I}, X_{II}, X_1 and $\mathbb{P}(W = 1) = p$;
- Z is independent of X_{I}, X_{II}, X_1 and $\mathbb{P}\ (Z = 1) = \frac{1}{2}$.

Both players learn the value of a Bernoulli random variable ξ, described as follows: the coin W is tossed by a referee; if $W = 1$, then the value of Y is revealed, otherwise the fair coin Z is tossed and the result of the toss is revealed. Thus

$$\xi = \begin{cases} Y & \text{if } W = 1, \\ Z & \text{if } W = 0, \end{cases} \tag{6}$$

i.e.

$$\{\xi = k\} = \{W = 1, Y = k\} \cup \{W = 0, Z = k\}, \quad k \in \{0, 1\}.$$

This mechanism is common knowledge, but the players don't know the outcome of W. They are only told the value of ξ. Observe that relevant information is given only when $W = 1$, which happens with probability p. If $p = 1$, then the players have an additional observation (exchangeable with X_1) before predicting X_1. If $p = 0$, then the additional information available (Y) is completely withheld. We may think of p as the intensity of the signal revealed. For each value of p we have a game, say G_p. We will consider different equilibria for G_p and compare their payoffs.

A strategy profile is described by a string of 8 actions. The first two are the actions taken by I_0 when $\xi = 0$ and $\xi = 1$, respectively, and so forth. Several equilibria emerge. Among these:

1. The strategy profile that emerges when each player plays a_1 if he thinks that $\{X_1 = 1\}$ is more likely than $\{X_1 = 0\}$, and a_0 otherwise, conditionally on the information available, including the signal ξ.
2. The same as above, but not taking into account the signal ξ.

For the remaining part of this section, we will assume that the parameters of the Beta prior law of Θ are $\alpha = 2.5$ and $\beta = 1$. We will show that:
(A) In case 1, the following strategy profiles emerge:

$$\overset{(I_0)}{a_0 a_1}\,\overset{(I_1)}{a_1 a_1}\,\overset{(II_0)}{a_0 a_1}\,\overset{(II_1)}{a_1 a_1} \quad \text{when } p > 11/19; \tag{7}$$

$$\overset{(I_0)}{a_1 a_1}\,\overset{(I_1)}{a_1 a_1}\,\overset{(II_0)}{a_1 a_1}\,\overset{(II_1)}{a_1 a_1} \quad \text{when } p \le 11/19. \tag{8}$$

These strategy profiles actually yield Nash equilibria, for the specified values of p.

(B) In case 2, instead, the strategy profile (8) emerges for all values of p.

(C) Equilibrium (8) is more favorable to the players than (7) when $p > 11/19$.

The message is that, no matter what the intensity of the signal is, the players achieve more by just ignoring it. Additional information correlated to the random variable to be predicted has no value for them.

3.3 Computations

Let us show, for example, what leads I_0 to the choices specified above when $\xi = 0$. We have

$$\begin{aligned}\mathbb{P}_{\mathrm{I}_0}(\xi = 0) &= \mathbb{P}_{\mathrm{I}_0}(W = 1, Y = 0) + \mathbb{P}_{\mathrm{I}_0}(W = 0, Z = 0)\\ &= p\frac{\beta+1}{\alpha+\beta+1} + (1-p)\frac{1}{2},\end{aligned}$$

and

$$\mathbb{P}_{\mathrm{I}_0}(X_{\mathrm{I}} = 0, \xi = 0) = p\frac{\beta+1}{\alpha+\beta+1}\,\frac{\beta+2}{\alpha+\beta+2} + (1-p)\frac{1}{2}\,\frac{\beta+1}{\alpha+\beta+1}$$

Hence,

$$\begin{aligned}\mathbb{P}_{\mathrm{I}_0}(X_{\mathrm{I}} = 0/\xi = 0) &= \frac{\dfrac{\beta+1}{\alpha+\beta+1}\left\{p\,\dfrac{\beta+2}{\alpha+\beta+2} + (1-p)\,\dfrac{1}{2}\right\}}{p\,\dfrac{\beta+1}{\alpha+\beta+1} + (1-p)\,\dfrac{1}{2}}\\ &= \frac{\dfrac{1}{22}p + \dfrac{1}{2}}{-\dfrac{1}{8}p + \dfrac{9}{8}} > \frac{1}{2} \Leftrightarrow p > \frac{11}{19}\end{aligned}$$

It is obvious that the strategy profile (8) yields a Nash equilibrium of G_p, for all values of p. We now show that the strategy profile (7) yields an equilibrium of G_p, for every $p > 11/19$.

Let $p > 11/19$. The payoff of I_0 in G_p if (7) is played is

$$\mathbb{P}_{\mathrm{I}_0}(\xi = 0, X_{\mathrm{II}} = 0, X_{\mathrm{I}} = 0) + \mathbb{P}_{\mathrm{I}_0}(\xi = 1, X_{\mathrm{I}} = 1) \tag{9}$$

$$\begin{aligned}&= \mathbb{P}_{\mathrm{I}_0}(W = 0, Z = 0, X_{\mathrm{II}} = 0, X_{\mathrm{I}} = 0)\\ &\quad + \mathbb{P}_{\mathrm{I}_0}(W = 1, Y = 0, X_{\mathrm{II}} = 0, X_{\mathrm{I}} = 0) + \mathbb{P}_{\mathrm{I}_0}(\xi = 1, X_{\mathrm{I}} = 1)\\ &= \frac{1-p}{2}\,\frac{\beta+1}{\alpha+\beta+1}\,\frac{\beta+2}{\alpha+\beta+2} + p\frac{\beta+1}{\alpha+\beta+1}\,\frac{\beta+2}{\alpha+\beta+2}\,\frac{\beta+3}{\alpha+\beta+3}\\ &\quad + \mathbb{P}_{\mathrm{I}_0}(\xi = 1, X_{\mathrm{I}} = 1)\\ &= \frac{4}{33} + \frac{4}{143}p + \mathbb{P}_{\mathrm{I}_0}(\xi = 1, X_{\mathrm{I}} = 1)\end{aligned}$$

whereas if she unilaterally deviates from (7) and plays $a_1 a_1$ (obviously she has no interest in playing $a_1 a_0$ or $a_0 a_0$) her payoff becomes

$$\mathbb{P}_{I_0}(\xi = 0, X_{II} = 1, X_I = 1) + \mathbb{P}_{I_0}(\xi = 1, X_I = 1) \tag{10}$$

$$= \frac{1-p}{2}\frac{\alpha}{\alpha+\beta+1}\frac{\alpha+1}{\alpha+\beta+2} + p\frac{\beta+1}{\alpha+\beta+1}\frac{\alpha}{\alpha+\beta+2}\frac{\alpha+1}{\alpha+\beta+3}$$

$$+ \mathbb{P}_{I_0}(\xi = 1, X_I = 1)$$

$$= \frac{35}{198} - \frac{175}{2574}p + \mathbb{P}_{I_0}(\xi = 1, X_I = 1)$$

Hence, for $p > 11/19$ and with the specified values of α and β, (9) is greater than (10). Similar computations show that player I_1 always prefers to stick to (7) rather than unilaterally deviating and playing $a_0 a_1$, no matter what the value of p is.

Thus, for $p > 11/19$, both (7) and (8) yield an equilibrium of G_p. We now want to compare their payoff vectors. It is clear that it is actually enough to compare the payoffs to I_0. His payoff if (8) is played is

$$\mathbb{P}_{I_0}(X_I = 1) = \frac{\alpha}{\alpha+\beta+1} = \frac{5}{9}$$

In order to compare this value with (9), observe that

$$\frac{5}{9} = \mathbb{P}_{I_0}(X_1 = 1) = \mathbb{P}_{I_0}(\xi = 0, X_1 = 1) + \mathbb{P}_{I_0}(\xi = 1, X_1 = 1)$$

and

$$\mathbb{P}_{I_0}(\xi = 0, X_1 = 1) = \mathbb{P}_{I_0}(W = 0, Z = 0, X_1 = 1)$$

$$+ \mathbb{P}_{I_0}(W = 1, Y = 0, X_1 = 1)$$

$$= \frac{1-p}{2}\frac{\alpha}{\alpha+\beta+1} + p\frac{\beta+1}{\alpha+\beta+1}\frac{\alpha}{\alpha+\beta+2}$$

$$= \frac{5}{18} - \frac{5}{66}p > \frac{4}{33} + \frac{4}{143}p \;\; \forall p \in \left(\frac{11}{19}, 1\right]$$

Hence, we see that the equilibrium (8) is the only one emerging according to the indicated criteria when $p \leq 11/19$, and is more favourable to the players than (7) when $p > 11/19$.

References

Avenhaus, R., Von Stengel, B. and Zamir, S. (1995). Inspection games, preprint.

Bassan, B., Scarsini, M. and Zamir, S. (1998). Uniqueness of Pareto optima, coordination games and positive value of information, preprint.

Bernardo, J. and Smith, A. (1994). Bayesian Theory, John Wiley & Sons, Chichester.

Donoho, D.L. and Johnstone, I.M. (1996). Neo-classical, minimax problems, thresholding and adaptive function estimation, *Bernoulli* **2**, 39–62.

Donoho, D.L. and Johnstone, I.M. (1998). Minimax estimation via wavelet shrinkage. *Annals of Statistics* **26**, 879–921.

Donoho, D.L. and Johnstone, I.M. (1994). Minimax risk over l_p-balls for l_q-error, *Probability Theory and Related Fields* **99**, 277–303.

Harsanyi, J.C. (1967/68). Games with incomplete information played by 'Bayesian' players, Parts I, II, and III, *Management Science* **14**, 159–182, 320–334, and 486–502.

Schwarz, G. (1994). Game theory and statistics, in *Handbook of Game Theory*, Volume 2, Edited by R.J. Aumann and S. Hart, Elsevier.