Universidade Federal de Minas Gerais
Instituto de Ciências Exatas
Departamento de Estatística

On the Three Prisoners Paradox

R. H. Loschi, S. Wechsler, P. L. Iglesias

Relatório Técnico

RTP-04/2001

Relatório Técnico
Série Pesquisa

On the Three Prisoners Paradox

Rosangela H. Loschi

Departamento de Estatística, Universidade Federal de Minas Gerais, Caixa Postal 702

CEP 30123-17, Belo Horizonte, Brasil - E-mail: loschi@est.ufmg.br

Sergio Wechsler

Instituto de Matemática e Estatística - Universidade de São Paulo, Caixa Postal 66281

São Paulo, CEP 05315-970 Brasil - E-mail: sw@ime.usp.br

Pilar L. Iglesias

Departamento de Estadística, Pontificia Universidad Católica de Chile, Casilla 306

Santiago 22, Chile - E-mail: pliz@mat.puc.cl

April 26, 2001

Abstract

The three prisoner paradox (e.g.,Jeffrey, 1992) has been solved by assuming that the receipt of relevant information is planned from the beginning. This assumption allows the use of Bayesian conditioning. This paper presents alternative explanations for the paradox based on other probability updating procedures such as superconditioning and Jeffrey's rule for updating. The situations of unpredictability of receipt of information is then also considered. The formulation of the paradox in this temporal setting brings new insight to the problem.

Key words: Unpredictability, Bayesian conditioning, superconditioning, Jeffrey's rule, Bayesian statistics.

1 Introduction

The three prisoners paradox (Jeffrey, 1992) is a good way to explain some different procedures for the updating of probabilities and on the other hand these procedures provide new insight in the paradox.

The three prisoners paradox is the following problem described by Jeffrey(1992), pp. 122:

The three prisoners (A, B, C). Two are to be shot and the other freed; none is to know his fate until the morning. Prisoner A asks the warder to confide the name of one other than himself who will be shot, explaining that as there must be at least one, the warder won't be giving away anything relevant to A's own case. The warder agrees, and tells him that B will be shot. This cheers A up a little, by making his judgement probability for being freed rise from 1/3 *to* 1/2*. But that's silly: A knew already that one of the others would be shot, and (as he told the warder) he is no wiser about his own fate for knowing the name of some other victim.* (Jeffrey (1992) - pp.122)

Notice that the "paradox" arises because, as prisoner A tells the warder, the information given about the other two prisoners does not apprise prisoner A of his own condition. Thus, the opinion of A about the event "A will live" ought to be the same after the receipt of the information provided by the warder, i. e., the posterior probability of this event should also be 1/3.

Usually, it is assumed that prisoner A constructs his posterior distribution using the Bayesian Conditioning (Jeffrey(1992), Howson and Urbach(1993), Howson(1996)) and that A calculates this distribution erroneously as he establishes an inadequate sample space to the experiment he is conducting - that is, A considers the sample space $\Omega_1 = \{A, B, C\}$, where X represents "prisoner X will live", $X = A, B, C$, and not the suitable space to the situation described in the paradox, say, $\Omega_2 = \{Ab, Ac, Bc, Cb\}$, where Xy represents the event "prisoner X will live and the warder informs that prisoner Y will die" (see Jeffrey (1992) for further details).

Other solutions for this problem can be find in Walley (1991) and Pereira (1995). Morgan et al. (1991) presents a valuable discussion about possible solutions to a similar problem. From Morgan et al.(1991)'s point-of-view, one concludes that the three prisoners paradox will only be solved if the strategy adopted by the warder is known, or if the Bayesian conditioning were used.

The Bayesian conditioning is a well known and widely used procedure for the updating of probabilities. Even so, there are several practical situations that do not allows us to use the Bayesian conditioning to obtain the posterior distribution. In next section we can see some details about the fragility of this rule for revising probabilities. (Procedure for the updating of probabilities is to be understood as the procedure which allows he/she to construct his/her posterior distribution from his/her prior distribution).

In this paper we present alternative explanations to the three prisoners paradox using three procedures for the updating of probabilities - the Bayesian conditioning, the superconditioning (Diaconis and Zabell(1982)) and the Jeffrey's rule(Jeffrey(1965,1992)). It will be argued that the solution presented by A may be correct if we do not arbitrate, as it is usually done, that: the receipt of information about B and C is planned from the beginning; the information received always takes A to the certainty that B will die; and prisoner A judges that the warder has the same chance to reveal the names of B and C, in case A is the one to be freed. Therefore we will point out some limitations to the Bayesian conditioning as well as highlight the influence of the way by which the information is received on the construction of the posterior distribution. It will be assumed that the prior probabilities declared by prisoner A for the events "prisoner X will live", $X = A, B, C$, are the same.

In Section 2, we will briefly present the Bayesian conditioning and responses to that paradox using this procedure for the updating of probabilities. In Section 3, prisoner A will be allowed to update his probabilities using the superconditioning defined by Diaconis and Zabell (1992). Finally, in section 4, the three prisoners paradox will be discussed in the light of the Jeffrey's rule.

In this paper, we denote by P and P^* two probability measures defined on the measurable space $(\Omega, \mathcal{A})$, where Ω is a countable set, and interpret P and P^* as the prior and posterior opinions of prisoner A about events in $\mathcal{A}$, respectively.

2 The Calculus of Prisoner A and the Bayesian Conditioning

As is well known, the Bayesian conditioning is the procedure for the updating of probabilities which connects prior and posterior distributions using the Bayes' theorem.

Definition 2.1 *(Bayesian Conditioning)*

The posterior distribution P^* is obtained from the prior distribution P using the Bayesian Conditioning if

$$P^*(.) = P(.|E) \tag{1}$$

where E is an $\mathcal{A}$- measurable event such that $P(E) > 0$.

The following result provide conditions under which the Bayesian Conditioning can be an acceptable procedure to obtain the posterior distribution.

Theorem 2.1 *(Jeffrey, 1992)*

Let $E \in \mathcal{A}$ be an event such that $P(E) > 0$. Then, for every $A \in \mathcal{A}$

$$\text{(i)} \quad P^*(E) = 1 \text{ and} \tag{2}$$

$$\text{(ii)} \; P^*(A|E) = P(A|E) \tag{3}$$

if, and only if,

$$P^*(A) = P(A|E). \tag{4}$$

The proof follows from the properties of probability measures.

Notice that the Bayesian conditioning is a procedure that allows he/she to update his/her initial opinions using the Bayes' theorem. When he/she decides to use it though the conditions (i) and (ii) in the previous theorem must be satisfied. This result is called by Howson and Urbach(1993) **The Principle of Bayesian Conditionalisation** and the conditions (i) and (ii) are also known as **certainty** and **sufficiency** respectively. Using this terminology it can be stated that the Bayesian Conditioning is an acceptable procedure for the updating of probabilities if the information received makes him/her move from an initial state of uncertainty about the conditioning event E to the posterior certainty of its occurrence and if, beyond this, the partition $\{E, \bar{E}\}$, generated by the information received, contains all the relevant information to the construction of his/her posterior distribution - i.e., $\{E, \bar{E}\}$ is a sufficient partition to the family of probability distributions $\{P, P^*\}$.

If, from the beginning, prisoner A plans to ask the warder about the situation of the other two prisoners, the sample space which appropriately describes the experiment performed by A is the space Ω_2 defined in section 1. Thus, a priori we have that $P(Ab) = p = 1/3 - P(Ac)$ and $P(Bc) = P(Cb) = 1/3$, where $p \in (0, 1)$.

At the very moment A declares his prior distribution, he also reveals his probabilities, supposing that the event "the warder informs that B will die" occurs. Denote this event by E and notice that $E = Ab \cup Cb$. Supposing that E occurs, the prior probability of the event "A will live" is given by:

$$P(Ab \cup Ac|E) = \frac{P(Ab)}{P(E)} = \frac{3p}{3p+1}. \quad (5)$$

If prisoner A judges that the conditional probabilities in E declared before are kept after consulting the warder and if his new opinion about the event E is $P^*(E) = 1$, the Bayesian conditioning allows A to explicit his posterior opinion for the event "A will live" as the probability obtained in the expression (5), i.e., $P^*(A) = P(Ab \cup Ac|E) = \frac{3p}{3p+1}$. Notice that this posterior probability is in accordance with the initial expectation about the posterior distribution that should be declared by prisoner A only if $p = 1/6$. This choice of p shows that, for A, the warder has equal chances to tell the names of either B or C , in case A is the prisoner who will live. On the other hand, if A thinks the warder would never tell the name of C, in case A were the survivor, that is, if he declares $p = 1/3$, the posterior distribution provided in (5) would agree with the posterior distribution divulged by prisoner A. Other approach to a similar problem can be found in Morgan et al.(1991).

Conversely, suppose that prisoner A has not planned to ask the warder about the condition of prisoners B and C. In this case, the suitable sample space to describe the experiment performed by A is the space Ω_1 defined in section 1.

As it has already been assumed, consider that each prisoner has a 1/3 chance of being the survivor . Consequently, the prior probability for the event "A will live", supposing that B will die, is given by:

$$P(A|E) = \frac{P(A)}{P(A \cup C)} = 1/2, \tag{6}$$

where $E = A \cup C$.

After A states his opinion about his being the one who will survive, the warder tells him that B will die. When he hears it, A decides to re-assess his initial probabilities using the unexpected information given by the warder. If in possession of the information given by the warder, A judges that event E is certain and that all the conditional probabilities in this event are not modified, the posterior probability for the event "A will live" is given by the expression (6), confirming prisoner A's initial statement.

In both situations we assume that the information received is the same and the Bayesian conditioning is the procedure adopted to update probabilities. Yet only in some situations the posterior distribution for the event "A will live" coincides with what seemed intuitive and logical at first. Besides, if the information is received unexpectedly, prisoner A will always be right.

Notice that the way by which the information is received influences the posterior probability calculus for the event "A will live" as it interferes directly with the construction of the sample space appropriate for the problem. In case the receipt of information is not previously anticipated, the change in the value of the posterior probability declared by A is plainly justifiable (in this case there is a change in the expectation), and the assessment made by prisoner A is not contradictory.

3 Explaining the Paradox via Superconditioning

It should be remarked, than Bayesian conditioning is not completely general, but rather it is valid under some specific assumptions as shown in section 2 (see also Diaconis and Zabell, 1982). Another example illustrating the reduced flexibility of this procedure for the updating of probabilities is presented in the following approach to the three prisoners paradox.

Suppose that prisoner A has not planned to ask the warder for information - i.e., A considers the space $\Omega_1 = \{A, B, C\}$ defined in section 1. Admit that prisoner A's prior probabilities for the survival of each prisoner is 1/3.

After eliciting his prior distribution for the events of Ω_1 (and before declaring his posterior distribution), prisoner A decides to ask the warder which of the other two prisoners will be sentenced, alleging that this information does not tell anything about his own condition. By doing so, prisoner A performs an experiment whose possible results are in the sample space $\Omega_2 = \{Ab, Ac, Bc, Cb\}$, where Xy is as in section 2. Suppose that, before asking the warder, prisoner A specify the following probability measure Q about events of Ω_2:

$$Q(Ab) = q_1; \quad Q(Ac) = q_2; \quad Q(Bc) = q_3 \ \text{ e } \ Q(Cb) = q_4,$$

where $q_i \in [0, 1]$, $\forall i$ and $\sum_i q_i = 1$.

How can the posterior distribution for the events of Ω_1 be determined? Notice that this situation has a little difference from that one described in section 2, where prisoner A has planned to ask for information before declaring his prior distribution. Here the Bayesian conditioning is not applicable.

Diaconis and Zabell (1982) present an answer to similar problems. The update procedure suggested by these authors (the superconditioning) is more general than the Bayesian conditioning and offers us an alternative way to explain the three prisoners paradox.

Definition 3.1 *(Superconditioning)*

The posterior distribution P^* can be obtained from the prior distribution P by Superconditioning if there exist a probability space $(\bar{\Omega}, \bar{\mathcal{A}}, Q)$ and a class of events $D = \{E_w \in \bar{\mathcal{A}}, w \in \Omega\}$, such that:

(i) E_w occurs if, and only if $\{w\}$ occurs, $\forall\ w \in \Omega$;

(ii) $Q(E_w) = P(\{w\})$, $\forall w \in \Omega$ and

(iii) $P^*(\{w\}) = Q(E_w|E)$ for every $w \in \Omega$ and for an event $E \in \bar{\mathcal{A}}$ such that $Q(E) > 0$.

Notice that, as for the Bayesian conditioning, the updating of probabilities via superconditioning is obtained multiplying the prior distribution by an appropiate likelihood function.

Define $E_A = Ab \cup Ac$, $E_B = Bc$, $E_c = Cb$ and $E = Ab \cup Cb$. Then we have that $Q(E_A) = q_1 + q_2 = 1/3$, $Q(E_B) = q_3 = 1/3 = q_4 = Q(E_C)$ and $Q(E) = q_1 + 1/3$. Thus, using the Superconditioning, the posterior probability for the event "A will live" is

$$\begin{aligned} P^*(A) &= Q(E_A|E) \qquad (7) \\ &= \frac{Q(Ab)}{Q(E)} = \frac{q_1}{q_1 + 1/3}. \end{aligned}$$

At first, suppose that prisoner A thinks the warder has an equal preference for both prisoners B and C , in such a way that $Q(Ab) = Q(Ac) = 1/6$ and $Q(Bc) = Q(Cb) = 1/3$. From (8) we have that $P^*(A) = \frac{1/6}{1/2} = 1/3$, what coincides with prisoner A's prior opinion about this event and corresponds to what was expected initially.

On the other hand, if A suspects that the warder will never tell that prisoner C will die, in case he is the survivor - what makes A declare $Q(Ac) = 0$ and $Q(Ab) = Q(Bc) = Q(Cb) = 1/3$ - the posterior probability of A being the survivor is $P^*(A) = 1/2$, what confirms the reason for prisoner A's excitement after talking with the warder.

If prisoner A adopts the superconditioning and has not uniform prior distribution on singleton events of Ω, A will only change his initial opinion to $1/2$ if there is any reason to judge that $Q(Ab) = Q(Cb)$. In case A considers $Q(Cb) = 2Q(Ab)$ we will have that $P(A) = P^*(A) = 1/3$ and in other circumstances none of the initial statements will be correct.

However, we must notice that the posterior distribution P^* is not always obtained from the prior distribution P by superconditionig. In the three prisoners paradox we can always use this updating procedure because the sample space is finite, as we can see in next theorem from Diaconis and Zabell (1982).

Theorem 3.1 *(Diaconis and Zabell, 1982)*

Let P and P^* be probability measures defined on $(\Omega, \mathcal{A})$, where Ω is a countable set. P^* is obtained from P by superconditioning if, and only if, there is a constant $B \geq 1$ such that

$$P^*(w) \leq BP(w), \quad \forall \quad w \in \Omega. \tag{8}$$

In next section, another way to explain the calculus done by prisoner A will be show by indruducing the Jeffrey's rule for the updating of probabilities.

4 The Jeffrey's Rule and the Three Prisoners Paradox

Other possible scenario we can use to explain the three prisoner paradox is described in the following. Suppose it is not prisoner A's intention to ask the warder for information about prisoners B and C, i.e., admit the sample space $\Omega_1 = \{A, B, C\}$. As before, admit that the prior probability distribution stated by A for the singleton events of Ω_1 are all 1/3.

If, after having declared his prior distribution about the events of Ω_1, prisoner A is unexpectedly informed that B will die, and if this information makes A change his opinion about the event $E = A \cup C$ arbitrarily, establishing that $P^*(E) = p^* < 1$, how can A construct his posterior distribution about the events of Ω_1? In these circumstances, the Bayesian conditioning and also the superconditioning do not offer a response to the problem.

As for the Bayesian conditioning, the events E and $\bar{E}$ define a partition of the sample space, but here it is not assumed that the information provided by the warder takes A to the certainty about the event E.

The Jeffrey's rule, presented in Jeffrey (1965, 1992), permits the construction of the posterior distribution in situations similar to the one we have just described.

Definition 4.1 *(Jeffrey's Rule)*

The posterior distribution P^* is obtained from the prior distribution P by Jeffrey's Rule if

$$P^*(.) = \sum_i P(.|E_i)P^*(E_i), \quad E_i \in \mathcal{E} \tag{9}$$

where $\mathcal{E} = \{E_i, i \in I\}$ is a partition of Ω, $P^*(E_i) \geq 0$ for every $i \in I$, $\sum_{i \in I} P^*(E_i) = 1$ and I is a index set.

The Jeffrey's rule also permits the arbitrary updating of the prior probabilities attributed to the elements of the partition. Its difference from the Bayesian Conditioning is that those arbitrary reassessments $P^*(E_i)$ may assume values smaller than 1 for every $i \in I$.

An inconvenience of the Jeffrey's rule is that it does not provide the procedures to obtain the posterior probabilities for the elements of the partition. To elicit those new probabilities can be such a psychologically complex task as stating the prior distribution.

Another inconvenience of the Jeffrey's rule is that in the way it is defined it does not provide a coherent posterior opinion in the sense defined by de Finetti(1931,1937) and Loschi(1992). A posterior probability measure will be obtained in this situations, if the conditions of the following theorem are satisfied.

Theorem 4.1 *Let $(\Omega, \mathcal{A})$ be a probability space where Ω is a countable set. Consider an $\mathcal{A}$-measurable event A and suppose that for every $E_i \in \mathcal{E}$, $P^*(E_i) > 0$. Then*

$$P^*(A) = \sum_i P(A|E_i)P^*(E_i) \tag{10}$$

if, and only if, for every $A \in \mathcal{A}$ and for every $E_i \in \mathcal{E}$

$$P(A|E_i) = P^*(A|E_i). \tag{11}$$

The proof of this theorem and a valuable discussion on Jeffrey's rule can be found in Jeffrey (1992). Some of its mathematical properties can be found in Diaconis and Zabell (1982) and Loschi, Iglesias and Arellano-Valle (2000) provide an application of this rule.

Let us admit that the condition (11) above is verified for prisoner A's prior and posterior opinions - i.e., for A the partition $\mathcal{E} = \{E, \bar{E}\}$, generated from the information provided by the warder, is sufficient to the family of probability measures $\{P, P^*\}$.

Since $P^*(E) = p^* < 1$, the posterior probability obtained by Jeffrey's Rule for the event "A will live" is:

$$\begin{aligned} P^*(A) &= P(A|E)p^* + P(A|\bar{E})(1 - p^*) \\ &= p^*/2. \end{aligned} \tag{12}$$

The only reason to prisoner A's excitement about having his chance of survival increased to 1/2 is when $p^* = 1$ - what would be the same as using the Bayesian Conditioning. The prior distribution of prisoner A remains unchanged afterwards if the information provided by the warder makes A less uncertain about the event E, but not totally convinced of its occurrence. In fact, A would have to declare $p^* = 2/3$ to have his prior distribution unchanged. In any other circumstances the posterior probability for the event "A will live" will be different from the prior probability and also different from the posterior probability established by prisoner A.

If, before stating his prior distribution, prisoner A intends to inquire the warder about his companions, the information that B will die causes the occurrence of the event $Ab \cup Cb$ of the sample space Ω_2. In this case, updating probabilities via Jeffrey's rule is equivalent to an updating by the Bayesian conditioning; thus the posterior probability of A being the survivor is given in (5). Notice that the Jeffrey's rule also does not apply to the problem considered by Morgan et al. (1991). In that problem, to use the Jeffrey's rule always corresponds to consider the Bayesian conditioning.

5 Conclusion

In this paper we present several scenarios to explain the three prisoner paradox and different approachs are considered to explain it. However, the problem is more general and often occurs in practical situations. Suppose, for example, we are predicting a time series with a truly dynamic model. Then, before, seeing more data, news broke that Irak invaded Kuwait. How are we going to make use of this relevant piece of information? In order to improve our predictions, this unpredictable information must be incorporated in the model. Notice that this real problem is very similar to the problem lived by prisoner A.

The three prisoners paradox, as well as any other problem which involves the updating of probabilities, does not have an unique solution (for the Three Prisoners Paradox we have seen that the number of solutions can be very large, and not only 1/2 and 1/3 as it is usually considered). One of the reasons is the lack of a normative rule for the choice of the procedure to be used for the updating of probabilities. In the absence of such a rule we may choose the procedure we judge the most adequate to the construction of our posterior distribution. We may even construct the posterior distribution by means of a completely arbitrary reassessment, that is, without using any mathematical formula.

Besides, the way by which the information is received influences our judgement. For all the situations discussed, the information was the same. How it was obtained as well as its influence on prisoner A's way of thinking generated distinct posterior distributions though. We do not know how prisoner A conducted his experiment, nor do we know which procedure he used to construct his posterior distribution. Therefore we cannot affirm that A made a mistake declaring 1/2 as the posterior probability of his being the survivor, because even arbitrary probability reassessments are permitted.

6 Acknowledgements

The authors like to thank Professor Luis Raul Pericchi (Univ. Simon Boliver-Venezuela) for his valuable suggestions. R.H. Loschi received partial financial support from CAPES, Brazil and P.L. Iglesias by FONDECYT, Chile (grant 8000004).

7 Bibliography

De Finetti, B.(1937). La prévision: ses lois logiqu

De finetti, B.(1931). Sul Significato Soggettivo della Probabilità. *Fundamenta Mathematicae* 17, 298-329.

Diaconis, P. and S. L. Zabell(1982). Updating Subjective Probability. *J. American Stat. Assoc.* 77, 822-830, No. 380.

Hacking, I.(1967). Slightly More Realistic Personal Probability. *Phil. Scien.* 34, 311-325.

Howson, C. and P. Urbach(1993). *Scientific Reasoning: The Bayesian Approach.* Second edition. Open Court, Chicago.

Howson, C.(1996). Bayesian Rules of Updating. *Erkenntnis* 45,195-208, No.2 and 3.

Jeffrey, R.(1992). *Probability and the Art of Judgment.* Cambridge University Press, Cambridge.

Jeffrey, R.(1965). *The Logic of Decision.* McGraw-Hill, New York.

Loschi, R.H., P.L. Iglesias and S. Wechsler(1998). Unpredictability and Probability Updating. To appear in *Rev. Soc. Chilena Estad.*

Loschi, R.H., P.L. Iglesias and R. B. Arellano-Valle (2000). Forecast of Returns in Stock Markets by Jeffrey's Rule. *Tech. Report RTP-08/2000* - Departamento de Estatística - Universidade Federal de Minas Gerais.

Loschi, R. H. (1992). *Coerência, probabilidad e Calibração.* Ms. Thesis. Universidade de São Paulo. (in portuguese).

Morgan, J.P., N.R. Chaganty, R.C. Dahiya and M.J. Doviak (1991). Let's Make a Deal: The Player's Dilemma. *American Stat.* 45, 284-289, No. 4.

Pereira, C.A.B.(1995). Estatística e Informação. *Boleletim da ABE*, Ano XI, No.30. (in portuguese)

Walley, P.(1991). *Statistical Reasoning with Imprecise Probabilities.* Chapmam and Hall, New York.