

ON RATIONALITY IN THE TRAVELER'S DILEMMA

RODRIGO MORO

Universidad Nacional del Sur

Instituto de Investigaciones Económicas y Sociales del Sur

(IIESS-CONICET)

rmoro@uns.edu.ar

MARCELO AUDAY

Universidad Nacional del Sur

Instituto de Investigaciones Económicas y Sociales del Sur

(IIESS-CONICET)

ccauday@criba.edu.ar

FERNANDO TOHMÉ

Universidad Nacional del Sur

Instituto de Matemática de Bahía Blanca

(INMABB-CONICET)

ftohme@criba.edu.ar

SUMMARY: Kaushik Basu presents the Traveler's Dilemma (TD) as a challenge to Game Theory. This challenge has been experimentally investigated. When faced with Basu's version of the TD, participants (including experts in game theory) behave in the way Basu suggests. However, a little change in the game turns out to reverse participants' choices. The question is, then, whether it is possible to provide an account of the main empirical findings as consequences of rational choices (i.e., to *rationalize* them). There are several proposals in the literature but none of them provides a satisfactory account for why experts in game theory playing against each other usually reject the only undominated option of the TD. The goal of this article is to suggest an alternative proposal that fixes this problem.

KEY WORDS: game theory, experiments, experts' choices, rationalization, domination

RESUMEN: Kaushik Basu presenta el dilema del viajero (DV) como un desafío a la teoría de juegos. El desafío ha sido investigado experimentalmente. Al enfrentarse a la versión de Basu del DV o similares, los participantes (incluyendo a expertos en teoría de juegos) se comportan como sugiere Basu. Sin embargo, un pequeño cambio en el juego tiene como consecuencia revertir las tendencias de elección. La cuestión es, entonces, si es posible brindar una explicación de los principales hallazgos empíricos como consecuencias de decisiones racionales (es decir, *racionalizarlos*). Hay varias propuestas en la literatura pero ninguna provee una explicación satisfactoria de por qué expertos en teoría de juegos jugando entre sí usualmente rechazan la única estrategia no dominada del DV. El objetivo de este artículo es sugerir una propuesta alternativa que solucione este problema.

PALABRAS CLAVE: teoría de juegos, experimentos, elecciones de expertos, raciona- lización, dominación

1. Introduction

In his paper of 1994, Kaushik Basu presents the Traveler's Dilemma (TD, henceforth) as a challenge to Game Theory. He illustrates the game with the story of two travelers who, returning from a remote island, find that the airline has damaged identical antiques that each had purchased. Thus, the airline manager, who does not know the price of the broken objects, designs the TD to determine how to compensate the travelers. An abstract version of the one-shot dilemma is the following:

You and another person face the following situation:

You have to choose a single number between 2 and 100.

The other participant has to do the same, but communication is not allowed.

In the case you and the other person choose the same number, each of you receives exactly that amount.

In the case you and the other person choose different numbers, payments are made considering only the *lower* of the two numbers chosen. The player who chooses the lower amount receives that *lower* amount *plus* \$2. The player who chooses the higher amount receives the *lower* amount *minus* \$2.

What number do you choose?

The standard analysis of this situation in game theory goes as follows. The three key assumptions are: 1) Agents are rational, in the sense of being utility maximizers; 2) The utility function of an agent is exclusively determined by the amount of money such agent makes; and 3) Rationality is common knowledge (everybody is rational, and everybody knows that everybody is rational, and everybody knows that everybody knows that everybody is rational, and so on *ad infinitum*). Let us go back to the game. Notice that strategy 100 is weakly dominated by 99, so no rational player would choose it. But this is common knowledge, so strategy 100 can be eliminated. If 100 is eliminated, 99 is weakly dominated by 98, and so on down to 2 (the only strategy that is not dominated). Thus, this game has only one Nash equilibrium, which is also the only rationalizable equilibrium, namely, the pair of strategies in the lowest end (i.e., {2, 2} in the version given above).

However, Basu (2007) argues as follows:

yet most people pick 100 or a number close to 100 —both those who have not thought through the logic and those who fully understand that they are deviating markedly from the “rational” choice. Furthermore, players reap a greater reward by not adhering to reason in this way. Thus, there is something rational about choosing not to be rational when playing Traveler’s Dilemma. (pp. 90–91)

Thus, according to Basu, the solution by game theory would go against our intuition on what is rational to do in the TD. This challenge is important, since game theory is almost universally taken as the normative standard for actions and decision in interactive situations.

Actually, there is empirical support for Basu’s challenge. Most participants in TD experiments choose high values, being the highest value of the range the typical modal response. Only a small percentage (between 20 and 30%) of participants choose the value prescribed by game theory. Rubinstein (2007) conducted the most extensive experiment with a sample of more than 1000 participants from several universities in different countries, systematically founding this pattern of results.

The reader may object that experimental studies (usually involving undergraduate students) may not be a reliable source of support. This is a perfectly reasonable objection. However, there is a special piece of evidence hard to dismiss. Becker et al. (2005) ran an experiment whose participants were members of the Game Theory Society, so they were (presumably) all experts in game theory, knowing that they were playing with other experts. The invitation packet even included Basu’s 1994 paper. The result was astonishing. Only around 10% chose the option prescribed by game theory, 20% chose the maximum value and the remaining 70% chose values in the 80s and 90s. Thus, most experts in game theory chose values that were consistent with Basu’s suggestion.

There is, however, an important empirical result that goes in the opposite direction. The literature has also taken into account another version of the TD: same rules as the TD but with a penalty and prize very high, say \$80. What is, then, our intuition in such a context? Is it still reasonable to play a high number? It is easy to check that if one does not play the Nash equilibrium strategy, one can end up with a \$78 debt. Thus, choosing a high number becomes extremely risky. Furthermore, the intuition is no longer obvious that one should choose a high number. What did experimenters find? Capra et al. (1999), Goeree and Holt (2001), Basu et al. (2011), and

Morone et al. (2014) showed that when the prize-penalty parameter is high (e.g., \$80 in a range of 80–200), participants tend to play low numbers with the mode in the lowest option. Moro et al. (2012) went a step further: if either part of the parameter is highly increased (either the prize or the penalty) leaving the other untouched, participants also tend to play low numbers. Clearly, we are in presence of the same game, or at least, a game with exactly the same structure. However, when the prize-penalty parameter is high, participants tend to play in the opposite direction, that is, in the direction that the standard analysis from game theory indicates.

The question is, then, whether it is possible to provide a rationalization of the observed behaviors, that is, to account them as consequences of rational choices. The literature has offered several accounts (we will review them in the next section) but they all involved weakening the agents' belief in the rationality of the other player. Thus, if an agent believes that the other player is not rational, the process of eliminating dominated strategies stops at the very beginning (except for the highest option). We will argue that these accounts are perfectly reasonable explanations for choices in typical contexts. However, they do not satisfactorily explain the following phenomenon: in the context of the original TD (i.e., with a low prize-penalty parameter) experts in game theory playing against each other tend to avoid the only undominated option of the game. In other words, these accounts provide no explicit reason for rejecting the option prescribed by game theory. Thus, rational players knowing that they are playing with peers should keep choosing it. The goal of this paper is to suggest an alternative account that fixes this problem. In order to do so, we will need to redefine "rational", since it cannot obviously have the same technical meaning as in game theory. Nevertheless, we will use a notion of rationality that underlies such a technical meaning, namely, the instrumental notion of rationality.

The paper is structured as follows: section 2 gives a brief review of the rationality accounts found in the literature, section 3 sets forth our own proposal, finally, section 4 presents our conclusions.

2. *Some Rational Accounts of Choices in the TD*

Becker et al. (2005) attempt to provide a rationality account of the main finding they report with game theory experts. They propose a model where the TD is analyzed as a game of incomplete information where there are three types of players: the irrational cooperative type, formed by agents who always play the highest option (i.e., 100 in the

original setup); the irrational uncooperative type, formed by agents who always play the lowest option (i.e., 2 in the original setup); and the rational type, formed by agents who try to give the best response to the opponent's choice. As a result, the existence of the irrational cooperative type leads the rational players to choose high strategies.

It is obvious that this proposal is a *partial* rationalization of behavior since some choices are classified as unavoidably irrational. As the authors claim,

Not using a strictly dominated strategy is among the most direct consequences of assuming rationality. So while one might for example use the lack of common knowledge of rationality to construct models in which rational players use strategies other than the Nash equilibrium strategy 2, there is no hope to explain $s = 100$ being played. (Becker et al., 2005, p. 8)

Thus, the authors refuse to provide a rationality account of those players choosing the highest option. Nevertheless, again, they are assumed to be game theory experts, so it is reasonable to demand a reason for such an anomalous behavior.

In order to provide another rationalization of behavior in the TD, Halpern and Pass (2012) propose a new solution concept to replace the Nash equilibrium: *repeated regret minimization*. Let us establish first the meaning of "regret" used by the authors. In informal terms, the value of regret is given by *what the agent missed obtaining* by choosing a certain option. For example, take again the original version of the TD and suppose agent 1 chose 90 and player 2 chose 95, obtaining \$92 and \$88, respectively. The value of regret of player 1 is 4, because he could have obtained \$96 (if he had chosen 94) instead of \$92. The idea is, then, that agents are uncertain about what the other agents would do and consequently choose an option that minimizes their regret. More precisely, agents would try to minimize the maximum amount of regret (i.e., they follow a minimax rule). It is easy to check that if a player chooses a value in the interval 96–100, he would have at most a regret of 3 for all possible choices of the other player. But for options lower than 96 the value of regret may be higher. So, if an agent wants to minimize his regret, he should choose in the interval 96–100. This result seems to match, at the qualitative level, the experimental results with a low prize-penalty parameter mentioned above (Capra et al. 1999; Rubinstein 2007; Basu et al. 2011). Additionally, if the prize-parameter is high (50 or higher), the only option that minimizes regret is the minimal value, 2, which also

matches, at the qualitative level, the experimental findings in TDs with a high prize-penalty parameter (Capra et al. 1999; Basu et al. 2011).

Going back to the analysis of the original TD, the result is, then, that options 2–95 do not minimize regret (notice that strategy 2 actually *maximizes* regret). Halpern and Pass propose, then, to delete those non-minimizing-regret strategies and repeat the process with the remaining interval, 96–100. The result is that option 97 is the one that minimizes regret, so it is the one that should be chosen.

This proposal is extremely interesting. However, it is easy to see that it cannot be assumed that it is common knowledge that both agents pursue the goal of minimizing regret. If player 1 predicts that player 2 would arrive to the conclusion that 97 is the option to choose, the option that minimizes the regret of player 1 is 96 rather than 97. But then the option that minimizes the regret of player 2 is 95, and so on to finally reach the situation in which both players choose strategy 2.

Halpern and Pass (2012) recognize this problem, so they postulate an assumption weaker than common knowledge of rationality (CKR): agents now assign decreasing degrees of probability to higher iterations of beliefs in rationality, so that the belief in the validity of CKR is extremely low. As a consequence, agents now believe with an overwhelming probability that the other agent will choose an arbitrary option, so 97 is still the rational option to pick.

Bach and Perea (2014) offer yet another possible rationalization of the empirical findings in the TD and similar games. Basically, they postulate a model in which players assign a positive probability to irrational choices of the other player. More specifically, agents have *utility-proportional beliefs*: opponent's better choices receive higher probability than inferior ones. In other words, probabilities of opponent's choices are proportional to the utilities the respective choices generate. As a result, in the context of the original TD, the model predicts that players will play around 96.

Thus, all the proposals found in the literature are based on weakening the agents' belief in the rationality of the other player. Again, these are perfectly reasonable accounts for choices in typical contexts. However, can they be used to explain the choices of game theory experts playing against each other? It is reasonable to assume that all of them can easily determine which one is the only undominated option in the game. So why don't they choose it? The above proposals are completely silent about this point. Let us, then, present our suggestion to fix this problem.

3. Our Proposal

As we mention above, we will use the notion of instrumental rationality, that is, the notion that commands to use all the available means to reach the goal or goals one has in mind. The reason behind our choice is twofold. First, we think that such a notion is behind Basu's criticism, since he argues that the response from game theory is unacceptable for agents who "play ruthlessly to try to make as much money as possible" (Basu 2007, p. 94). Second, the instrumental notion of rationality is typically used in the context of actions and decisions. As usual, agents may differ in the goals they pursue, so what is rational may also vary accordingly. Thus, we will organize this section by considering different goals an agent may pursue when facing the TD.

3.1. The Goal of Getting a Payoff as High as Possible

Let us start by considering that the goal each agent pursues is *to get a payoff as high as possible*.¹ Let us also assume that this is common knowledge. The problem is that agents do not know what the other player is going to do. If they did, it would be very easy for them to find the option that gives them the highest payoff. Nevertheless, some options may have certain features that make it impossible to satisfy the proposed goal with them. Thus, agents should discard them. Let us begin by recognizing that standard game theory already does most of the job for us. The key concept is domination. Notice that to say that option x *strictly dominates* option y means that, no matter what the other player chooses, option x always gives you a higher payoff than option y . So, if an agent wants to get a payoff as high as possible, he or she should avoid strictly dominated options since it is possible to obtain higher payoffs with other options. Something similar can be argued about *weakly dominated* options: An option x *weakly dominates* option y if, no matter what the other player chooses, option x always gives you either a higher payoff (in, at least, one case) or the same payoff as option y . Thus, depending on what option the other player chooses, you may be indifferent between choosing x and choosing y . But in all the cases that x and y give you different payoffs, option x gives you a higher payoff than y . Thus, if one does not know what the other agent is going to choose, one should also

¹ We use this formulation of the goal to make it clear that it should not be confused with the concept of payoff maximization, which is linked to the technical meaning of rationality that leads to the paradoxical result in the TD. The difference between the two goals will become clear at the end of this section.

avoid weakly dominated options. Furthermore, if one knows that the other player pursues the same goal, one knows that the other player will also avoid both strictly and weakly dominated options. This is the reason to eliminate (or at least, to leave aside) options 3–100 in the original version of the TD. So, is it, then, rational to pick option 2, the lowest value of the range? There seems to be a general agreement (again, even for game theory experts playing the game) that this result goes against intuition, but it is necessary to show that it also fails to meet the goal of getting a payoff as high as possible.

In order to justify this claim, we need to take an intermediary step by considering the concept of *strong Pareto domination*. An outcome of a given game is strongly Pareto-dominated when there exists, at least, an alternative outcome where each player gets a higher payoff. For example, considering the original TD, compare the outcome where each player obtains \$40 (by simultaneously choosing 40) and the outcome where each obtains \$50 (by simultaneously choosing 50). We will say, then, that the outcome (\$40–\$40) is strongly Pareto-dominated because there is, at least, another one (i.e., \$50–\$50) where each player earns more money. Let us now consider this type of outcome from the perspective of an agent who wants to get a payoff as high as possible. It is clear that such an agent should try to avoid strongly Pareto-dominated outcomes, simply because he or she can get higher payoffs with other outcomes. Notice also that the other player, having the very same goal in mind, also wants to avoid such Pareto-dominated outcomes. These facts suggest that it is possible, in principle, to reach an implicit agreement to avoid this type of outcomes. It is worth pointing out that this is independent of any pro-social consideration. In other words, even completely selfish players should try to avoid Pareto-dominated results.

The problem is, however, that players do not choose between outcomes but between strategies. Furthermore, depending on what the other player chooses, the same strategy may lead to a Pareto-dominated outcome and to an outcome with the highest payoff of the game (e.g., defecting in the Prisoner's Dilemma). These are reasonable objections, but consider a strategy with the following feature: it *always* leads (*no matter what the other player chooses*) to strongly Pareto-dominated outcomes. Let us call these strategies “derivatively dominated”. The reason for such a name is that, although they may not be dominated per se, they always result in strongly Pareto-dominated outcomes.

It seems clear that in order to avoid Pareto-dominated outcomes, agents should avoid derivatively dominated strategies because this

type of strategy always leads to that type of outcome. As a consequence, if agents want to get a payoff as high as possible, they should avoid derivatively dominated strategies.

Finally, we have set up the elements to explain why strategy 2 in the original TD is not *the* rational option of the game: such an option is a derivatively dominated strategy. Let us show this in detail. If the other player chooses a strategy in the 3–100 range, our agent gets \$4 and the other player gets nothing. For all those combinations, there is at least an alternative combination of strategies, say (100–100) or (100–99) or (99–100) where each agent earns more money. If the other player also chooses strategy 2, both players get \$2. Again, such a combination is dominated by the same profiles. Thus, strategy 2 *always* leads to a result that is strongly Pareto-dominated by an alternative profile. Therefore, rational players trying to get a profit as high as possible would avoid it. Game theory does not take into account this alternative way of conceiving domination and this may be the reason behind the problematic solution to the TD.

It is easy to check that options in the range 2–97 are derivatively dominated. Therefore, they should also be eliminated (or, at least, left aside). Recall that strategies 3–100 should be eliminated or left aside for being dominated (100) or iteratively dominated (3–99) in the usual sense of the term. Unfortunately, the consequence is that there is no remaining option. In other words, there is no stable combination of strategies: For each profile of strategies, one or both players have an incentive to move away from it.

To sum up, the result of considering the goal of getting a payoff as high as possible for two rational players is that there is no rational choice in the TD because all the possible courses of action are in some way or another dominated (dominated, iteratively dominated or derivatively dominated).

So we are back to square 1 and we still have to make a decision. This leads to consider a different (although closely related) goal.

3.2. The Goal of Getting a High Payoff

Basu (1994) suggests as a solution to represent the problem in terms of ill-defined categories (i.e., to choose a high/low number). Thus, if one represents the original problem in such a way, the matrix of the game is as table on page 65.

Here the option of choosing a low number is both dominated in the usual sense and derivatively dominated, so it should not be chosen. Thus, by representing the situation in those terms, the rational

response becomes clear: each player should choose a high value, so both of them end up with a high payoff.

We think that this is a very reasonable proposal, at least for the original version of the TD or similar ones with a low prize-penalty parameter (we next analyze the case of a TD with a high prize-penalty parameter). However, notice that this amounts to an implicit change in the goal the agents are pursuing. For this proposal to work, agents should not longer try to get a payoff as high as possible. In order to see this, notice that the highest value of the range should surely count as a high value. But as we mentioned above, this value is dominated by the next lower value (i.e., 99 in the original version), so it can never be a maximizing strategy. Nevertheless, if the goal is just to obtain a high payoff, Basu's proposal provides a reasonable suggestion. Notice, however, that it has an important shortcoming and thus must be nuanced: it only works for agents representing the game by using *ill-defined* categories (Basu was the first to point out this restriction, see Basu 1994, pp. 394–395). In other words, if the category *high number* is taken to refer to a precise set of values, one starts analyzing the possibility of choosing one unit less than the lowest value of such a range and then, one unit less than such a value and so on. Following this line of analysis, the paradox arises again. Thus, it is not sufficient for both players to have the goal of obtaining a high payoff to reach a solution. They should also represent the game by using the suggested ill-defined categories and, of course, this must be common knowledge.

Before we go to the case of a high prize-penalty, it is interesting to consider the relationship between the two goals analyzed so far. It is possible to establish a hierarchy between them: the first goal to pursue is to get a payoff as high as possible but, if there is no rational option, one can go for the more modest goal of getting a high payoff, which is different but clearly goes in the same direction as the main goal.

Notice that it is not possible to apply the strategy of using ill-defined strategies when the prize-penalty parameter is high (for example, 80 in a 2–100 range). In order to see this point, consider the case where both players choose a high number. It is no longer possible to represent the result by using the previous ill-defined categories, since the result may turn out to be (high, high) or (high, low) or (low, high) depending on the exact numbers played (e.g., 90–90, 90–95 and 95–90, respectively). In other words, when the incentive parameter is very high, we lose the certainty that if both agents play

a high number, they will end up with a high payoff. Should we, then, abandon the game? Not necessarily, because it is possible to lower even more our expectations. If it is not possible to get a payoff as high as possible or get a high payoff, maybe it is possible at least to get the minimal payoff.

		Player II	
		High	Low
Player I	High	high, high	low, low
	Low	low, low	low, low

Table: Payoff matrix of the original Traveler's Dilemma presented in ill-defined categories

3.3. The Goal of Getting a Positive Payoff

In the TDs with a high prize-penalty used in the experimental literature (typically, 80 in a 80–200 range), by not choosing the minimal value, one faces the risk of ending up with nothing, which is clearly worse than getting some money out of the game (this claim is supported by the experimental studies where researchers ask about motivations in the TD, see Basu et al. 2011, and Brañas-Garza et al. 2011). Is it rational, then, to choose the minimal value when the prize-penalty is high? Again, the answer to this question depends on the goal the agents have in mind. If we want to get a payoff as high as possible, we are in the same situation as with the original version of the TD, where there is no stable combination. Indeed, strategy 80 is still derivatively dominated (e.g. with regard to profile 200–200). If the goal is to get a high payoff, it is no longer possible to represent the situation by using ill-defined categories as happens in TDs with a low prize-penalty parameter. Furthermore, it is clear that if both agents choose the minimal value, they are not getting a high payoff, relatively speaking. If the goal is, though, to get some money (i.e., a positive payoff), then the option of choosing the minimal value is the rational strategy, since it is the only way to make sure that one earns something in the game.

Again, the relationship between goals can be seen in the way suggested in the previous section, that is, as a hierarchy of goals.

The impossibility of finding a stable combination of strategies implies that, in turn, it is not possible to find a solution under the goals of either getting a payoff as high as possible or getting a high payoff. Thus, in this case it seems reasonable to try to get, at least, the minimal payoff of the game.

This concludes our analysis. Of course, it is possible to extend it to consider other goals the agents may have in mind as well as other beliefs about the other player, and their respective interactions. Nevertheless, we also think that this should be enough to present our initial view on the subject.

Before we present our conclusions, it is important to recognize that there are related theoretical and empirical investigations on the matter.

On the theoretical side, we should start by acknowledging that game theorists have incorporated many goals different from “monetary payoff maximization” in their analysis of games, usually to account for anomalous behavior. Thus, for example, there is a huge amount of literature on *social preferences*, where the utility function of an agent also depends on the payoff of the other player/s (for a review see Cabrales and Ponti 2011). Our analysis also relates to the study by Halpern and Rong (2010), since they propose a new equilibrium concept that also uses the notion of Pareto optimum. On the empirical side, there is a study of particular related interest, namely, Brañas-Garza et al. 2011. The authors obtained self-reported justifications from participants after playing the TD. The main conclusion from this study is that there exists a high heterogeneity of motivations while playing the TD, that is, different participants report different motivations (aspiration levels, competitive preferences, risk aversion, penalty aversion, etc.). More importantly, participants tend to make their decisions accordingly to their own motivations not only in the TD, but also in other games and tests (Basu et al. 2011 report somewhat similar results). Although this goes in line with our main argumentation, it is necessary to recognize that the wide variety of motivations of participants amply exceeds the limited scope of our analysis.

4. Conclusion

Kaushik Basu presents the Traveler’s Dilemma (TD) as a challenge to game theory. The empirical evidence seems to support Basu’s criticism. Of special importance is that even experts in game theory behave in the way Basu suggests. The question is, then, whether it is

possible to provide an account of such evidence in terms of rational decisions. We revised the proposals in the literature but none of them provided a satisfactory account for why experts in game theory playing against each other usually reject the only undominated option of the TD. We then tried to provide an alternative account that could fix this problem. In order to do so, we use the instrumental notion of rationality, which is typically applied to the assessment of actions and decisions. But the real key of our proposal is to introduce the concept of *derivatively dominated strategy*. Even if the lowest option of the range is undominated in the usual sense of the term, we show that it is derivatively dominated, that is, it always leads to strongly Pareto-dominated results. Every researcher seemed to accept that, in the context of the original TD, the lowest option is not a satisfactory response. However, there is no concrete proposal explaining why this was the case. Our proposal helps to articulate a concrete reason for such a rejection.

We finally recognize that our analysis can be extended to consider other goals and other beliefs about the other player. More importantly, we do not pretend to have offered the last word on the question of rationality in the TD. The main motivation of this paper is just to throw light on an issue that has not received the philosophical attention it deserves.²

REFERENCES

Bach, C. and A. Perea, 2014, "Utility Proportional Beliefs", *International Journal of Game Theory*, vol. 43, no. 4, pp. 881–902.

Basu, K., 2007, "The Traveler's Dilemma", *Scientific American*, vol. 296, no. 6, pp. 90–95.

_____, 1994, "The Traveler's Dilemma: Paradoxes of Rationality in Game Theory", *American Economic Review*, vol. 84, no. 2, pp. 391–395.

Basu, K., L. Becchetti, and L. Stanca, 2011, "Experiments with the Traveler's Dilemma: Welfare, Strategic Choice and Implicit Collusion", *Social Choice and Welfare*, vol. 37, no. 4, pp. 575–595.

Becker, T., M. Carter, and J. Naeve, 2005, "Experts Playing the Traveler's Dilemma", Discussion Paper 252, Institute for Economics, Hohenheim University, available at: <<https://www.uni-hohenheim.de/RePEc/hoh/papers/252.pdf>> [retrieved: 16/01/2018].

² We want to thank two anonymous reviewers for helpful comments and suggestions on a previous version of this paper.

Brañas-Garza, P., M. Espinosa, and P. Rey-Biel, 2011, “Traveler’s Types”, *Journal of Economic Behavior and Organization*, vol. 78, nos. 1–2, pp. 25–36.

Cabrales, A., and G. Ponti, 2011, “Preferencias Sociales”, in P. Brañas-Garza (ed.), *Economía experimental y del comportamiento*, Antoni Bosch, Barcelona.

Capra, C., J. Goeree, R. Gomez, and C. Holt, 1999, “Anomalous Behavior in a Traveler’s Dilemma?”, *American Economic Review*, vol. 89, no. 3, pp. 678–690.

Goeree, J.K., and C.A. Holt, 2001, “Ten Little Treasures of Game Theory and Ten Intuitive Contradictions”, *American Economic Review*, vol. 91, no. 5, pp. 1402–1422.

Halpern, J., and N. Rong, 2010, “Cooperative Equilibrium”, *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, vol. 1, International Foundation for Autonomous Agents and Multiagent Systems, Toronto, pp. 1465–1466; available at: <<http://ifaamas.org/Proceedings/aamas2010/pdf/02%20Extended%20Abstracts/Red/R-49.pdf>> [retrieved: 16/01/2018].

Halpern, J., and R. Pass, 2012, “Iterated Regret Minimization: A New Solution Concept”, *Games and Economic Behavior*, vol. 74, no. 1, pp. 184–207.

Moro, R., E. Freidin, and F. Contiggiani, 2012, “¿Zanahorias o palos? El efecto de premios y castigos en el dilema del viajero”, in L. Salvatico, M. Bozzoli and L. Pesenti (eds.), *Epistemología e Historia de la Ciencia: Selección de trabajos de las XXII Jornadas*, vol. 18, Universidad Nacional de Córdoba, Córdoba, pp. 405–411.

Morone, A., P. Morone, and A.R. Germani, 2014, “Individual and Group Behaviour in the Traveler’s Dilemma: An Experimental Study”, *Journal of Behavioral and Experimental Economics*, vol. 49, pp. 1–7.

Rubinstein, A., 2007, “Instinctive and Cognitive Reasoning: A Study of Response Times”, *The Economic Journal*, vol. 117, no. 523, pp. 1243–1259.

Received: July 1, 2016; revised: May 30, 2017; accepted: January 12, 2018.