האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

UNCERTAINTY IN THE TRAVELER'S DILEMMA

By

GILAD BAVLY

Discussion Paper # 595 January 2012

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: http://www.ratio.huji.ac.il/

Uncertainty in the Traveler's Dilemma*

Gilad Bavly†

January 10, 2012

Abstract

The paper analyzes a perturbation on the players' knowledge of the game in the traveler's dilemma, by introducing some uncertainty about the range of admissible actions. The ratio between changes in the outcomes and the size of perturbation is shown to grow exponentially in the range of the given game. This is consistent with the intuition that a wider range makes the outcome of the traveler's dilemma more paradoxical. We compare this with the growth of the elasticity index (Bavly (2011)) of this game.

1 Introduction

This paper deals with the effect that uncertainty can have on the analysis of the Traveler's Dilemma (Basu (1994)), a well-known "paradox of iterative reasoning."

The Traveler's Dilemma (TD) concerns two travelers, each claiming an integer sum between \$0 and \$100, as compensation for lost luggage. Upon

*This work is part of the author's Ph.D. thesis, done under the supervision of Professor Abraham Neyman. I am deeply grateful to Prof. Neyman for his kind and illuminating guidance. The research was supported in part by Israel Science Foundation grants 1123/06 and 1596/10.

†Institute of Mathematics and Center for the Study of Rationality and Interactive Decision Theory, Hebrew University, Jerusalem, Israel (*e-mail*: `gilbav@math.huji.ac.il`).

receiving these two claims, the airline considers the lowest of them to be the "truthful" one, and gives this amount to both travelers, but (if they claimed different amounts) with an extra bonus of \$2 to the low claimant, and a deduction of \$2 to the high claimant. Thus, the payoff of player i is ($i \neq j$, $i, j = 1, 2$):

$$h_i(x_i, x_j) = \begin{cases} x_i + 2 & \text{if } x_i < x_j \\ x_i & \text{if } x_i = x_j \\ x_j - 2 & \text{if } x_i > x_j \end{cases}$$

In addition to being the unique (correlated) equilibrium, the pair of actions $(0, 0)$ consists of the only actions that survive iterated deletion of strongly dominated strategies.[1][2] Yet, as a prediction, demanding only 0 may seem counterintuitive.[3]

Although the game here is a one-stage game, the iterative nature of the reasoning that leads to this perhaps paradoxical outcome may remind us of some well-known paradoxes of backward induction in multistage games (e.g., the centipede game, or the repeated prisoner's dilemma). We shall see more of this connection in the analysis to follow.

In Section 2 we analyze the effect that uncertainty about the range of possible demands can have on the outcome of this game. There is a close analogy between this kind of uncertainty, and the uncertainty about the number of stages in a repeated game of Neyman (1999). In line with the intuition that the wider the range, the more paradoxical the reasoning, we get that the size of the perturbation that is needed to reverse the analysis decreases in an order exponential in the range. In Section 4.1 we compare this concept of uncertainty about range with the elasticity index (Bavly (2011)),

[1]A strategy is *strongly dominated* if there exists another strategy, or a mixture of strategies, that receive a strictly higher payoff for any choice of the other players.

[2]it is also the unique rationalizable pair of actions, since in a 2-player game rationalizability is equivalent to iterated deletion of strongly dominated strategies.

[3]The more common story of TD is of demands between 2 and 100 (and this way no player can ever receive a negative payoff). We start from 0 for convenience, to make indices more readable later on.

which involves a more general kind of uncertainty.

2 Uncertainty about the Range

Let n denote the highest number that the players can demand (i.e., in our story n was 100), and consider some possible values of n (we denote the game with upper bound n as $(\mathrm{TD})_n$). If, for example, $n = 3$, then the $(0,0)$ outcome seems to make some sense. But most people will probably feel that for $n = 100$, "real players" will not play $(0,0)$. We maintain that, indeed, the real-life situation that we have in mind may not be adequately modelled by the game as given. What is more, this is the case even if the situation contains no incentives or possibilities, either implicit or explicit, for cooperation among the players.

The point is that the above iterative reasoning is dependent on the players having common knowledge of the game data and of their rationality. While this assumption may be unrealistic, it becomes more crucial to the analysis as the number n grows, because the larger n is, the smaller the "concessions" (namely, the smaller departures from this assumption) are that are needed to get the players to play something that is far from $(0,0)$. This fits well with the intuition that for large n, $(0,0)$ is not going to be played.

To demonstrate this, we focus here on one kind of departure from this assumption, and describe a system in which it is not common knowledge that the upper bound equals the number n. Knowing that the upper bound is n means knowing what is the highest demand that I can make, and knowing that demanding $n+1$ will "not be heard," or rather will just be perceived as demanding n.

The basic properties of the following construction can be described informally as follows. In the vast majority of states of the world, the players know the upper bound is n. Also in the vast majority of states they know that the other player knows that the range is n, and also when it comes to third level of knowing, etc., but not ad infinitum; i.e., there is no common knowledge of the range being n. Furthermore, the average perceived upper bound over all states is very close to n. And yet, the average actions they choose are far from $(0,0)$, in fact close to (n,n). And the players are playing a Bayesian equilibrium (i.e., there is common knowledge of rationality).

Let T denote the upper bound in the various states of the world. In order to account for various values of T, we take the set of actions available to each player to be not just integers between 0 and n, but all integers. Thus, in a state of the world where $T = k$, a player may in fact demand more than k, but this action is equivalent to demanding k. With this interpretation in mind, we can also formalize the original game (where the upper bound is deterministically n) with this infinite set of actions, where any demand larger than n is just equivalent to n. Then there will be $n+1$ equivalence classes of actions.

The uncertainty structure consists of:

(a) For each $i \in N$, a set S_i (i's types),

(b) For each $s_i \in S_i$, a pure action (i's choice) in A_i,

(c) A common prior p over $\Omega = \times_{i \in N} S_i$,

(d) For each $i \in N$ and $\omega \in \Omega$, a bound $T(\omega)$.

The payoff for player i at state ω is then $u_i(x_i, x_j)(\omega) = h_i(\min\{T(\omega), x_i\}, \min\{T(\omega), x_j\})$, where h_i is the payoff function depicted above.

A player is rational at state ω when her chosen action maximizes her payoff, given her information.

We will refer to such an uncertainty about T structure as T for short, and to the resulting game as $(\mathrm{TD})_T$.

We say that there is *k-mutual knowledge* of some event E at state ω, if for any list of players $(i_1, i_2, \ldots, i_l)$ whose length l is $\leq k$, it is true at state ω that player i_1 knows that i_2 knows .... that i_l knows E.

Theorem 2.1. *There exists a number K such that for any n and any $0 \leq x \leq n$, there is an uncertainty structure T with $T \geq n$ and $E(T-n) \leq \exp(-m(x))$, where $m(x) = \frac{n-x}{K}$, and with probability $\geq 1 - \exp(-m(x))$ there is $m(x)$-mutual knowledge that $\{T = n\}$, and there is an equilibrium of*

$(TD)_T$ with payoffs that are K-close to (x, x).

Corollary 2.2. *There exists a number C such that for any $0 \leq \lambda \leq 1$, $\varepsilon > 0$, and a sufficiently large n, there is an uncertainty structure T with $T \geq n$ and $E(T - n) \leq \exp(-\varepsilon n)$, and an equilibrium of $(TD)_T$ with payoffs that are $(C \cdot \varepsilon n)$-close to $(\lambda n, \lambda n)$.*

On the other hand, the exponential scale of the above theorems is indeed the most that can be done. The following theorem states that when the scale of uncertainty is smaller than that, we cannot deviate far from $(0, 0)$.

A little notation: $[x]_+ = max\{x, 0\}$.

Theorem 2.3. *There exists a number K such that for any $\varepsilon > 0$ there is N such that for every $n > N$, if $E([T - n]_+) < \exp(-Kn)$, then all equilibrium payoffs of $(TD)_T$ are ε-close to $(0, 0)$.*

With these "R-uncertainties," namely uncertainties about the range of actions, we may define the *R-elasticity* of the parameterized game as the maximal ratio between the expected change of equilibrium payoff and the expected change of the parameter[4] $(E|T - n|)$.

Corollary 2.4. *There exist numbers $B \geq A > 1$, such that the R-elasticity of $(TD)_n$ is larger than A^n and smaller than B^n.*

3 Proofs

Proof of Theorem 2.1. The structure has two parameters: a constant real number $q < 1$ sufficiently close to 1, and an integer $r = \lfloor x \rfloor$. Let the non-zero part of Ω be $\{\omega = (i, j) : i, j \geq 0,$ and $i - 3 \leq j \leq i + 3$ when $i < r,$ and $i - 3 \leq j \leq i + 4$ when $r \leq i \leq r + 3,$ and $i - 4 \leq j \leq i + 4$ when $r + 3 < i\}$, illustrated in Figure 1 (note that it is symmetric between i and j). At $\omega = (i, j)$ player 1 knows only what i is, and chooses the action i; and player

[4]In our case the parameter n governed the range, hence the name, but of course we may think of doing such a thing with other kinds of parameters.

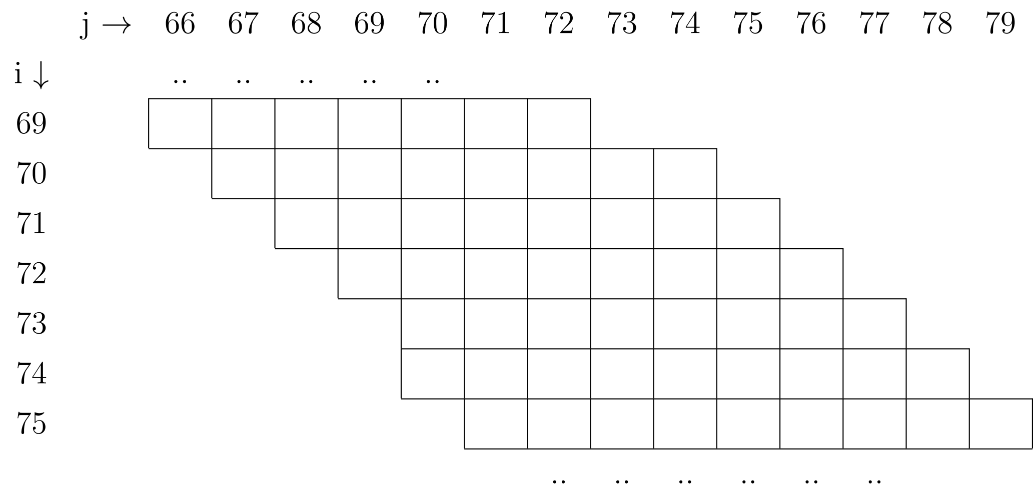

Figure 1: Part of the structure (with $r = 70$)

2 knows only what j is, and chooses the action j. $T(\omega) = \max\{n, i, j\}$, and let $p(\omega) = \alpha q^{\max\{|i-r|,|j-r|\}}$ (α is the appropriate normalization).

Note that if a player's conjecture is that the other player may choose any action $a \leq x \leq b$ (and $b - a \geq 4$), each x with equal probability, then her best replies are the fourth from above, $b-3$, and the fifth, $b-4$. One way to see this is to note that if your candidate for best reply is some x, and you check whether it is better to switch to $x+1$, then against anything below x the switch will not matter, against x the switch loses 2, against $x+1$ it loses 1, and against anything higher it gains 1.

If, instead of equal probability, the probability of the three upper actions (b, $b-1, b-2$) is slightly increased, then $b-3$ is the unique best reply; and if the probability of the three upper actions is slightly decreased, then $b-4$ is the unique best reply. It follows that when q is sufficiently close to 1 our choice of the prior p ensures that for any i, playing i is indeed the best reply against player 1's conjecture (for $i < r$, i is the fourth possible j, and for $i \geq r$ it is the fifth); and similarly for player 2. In other words, the players are rational at every state.

$E\,(T-n) = \sum_\omega (T(\omega) - n)\, p(\omega)$. Summing over the main diagonal (where

$i = j$) we have α times $\sum_i([i-n]_+)\, q^{|i-r|} = \sum_{k=1}^{\infty} k\, q^{n-r+k} = q^{n-r}\sum_{k=1}^{\infty} k\, q^k = q^{n-r}\frac{q}{(1-q)^2}$. Summing over the adjacent diagonals, where $i > j$, we get again the same sum $\sum_{k=1}^{\infty} k\, q^{n-r+k}$, and by symmetry between i and j, we should get the same over diagonals where $j > i$. So overall we get that $E\,(T-n)$ is proportional to q^{n-r}.

At a state $\omega = (i_0, j_0)$, player 1 knows that $j \leq i_0 + 4$, and therefore that 2 knows that $i \leq i_0 + 8$, etc. Therefore there is $\lfloor (n - \max\{i, j\})/4 \rfloor$-mutual knowledge that { T=n}. In particular, at the set $E = \{w = (i, j) : i, j \leq n - 4(n-r)/5\}$ there is $((n-r)/5)$-mutual knowledge that $\{T = n\}$. And since $q < 1$, most of the probability is concentrated near $\omega = (r, r)$. Specifically, to evaluate $p(\Omega \backslash E) = \sum_\omega p(\omega)$ let us first sum over the main diagonal: there we get α times $\sum_{i=n+(n-r)/5}^{\infty} q^{|i-r|} = \sum_{k=(n-r)/5}^{\infty} q^k = q^{(n-r)/5}\frac{1}{1-q}$. Summing over the other diagonals gives the same computation; therefore the probability is proportional to $q^{(n-r)/5}$.

Finally, we claim that the average value of i over all states is within a constant of r. There is a sort of symmetry around the point (r, r), albeit imperfect. One asymmetry is that the diagonal $\{i = j + 4\}$ only starts at $i \geq r$. The average of $(i-r)$ over this diagonal is α times $\sum_{i=r}^{\infty}(i-r)\, q^|i - r| = \sum_{k=0}^{\infty} k\, q^k = \frac{q}{(1-q)^2}$. A similar calculation applies for the asymmetry of the diagonal $\{i = j + 4\}$. And all the diagonals are infinite only on one side, but that only adds tails of the same sums. Now, at a state $w = (i, j)$ the sum of payoffs is $i + j$. Hence, the average sum of payoffs is the average of $i + j$, and by symmetry it is twice the average of i, and also by symmetry the average payoffs of 1 and 2 are equal. Therefore, the average payoff is within a constant of (r, r). □

Proof of Theorem 2.3. We first define a stochastic game Γ_T that is payoff-equivalent to the TD. The game starts at state s_1, where each player i has two available actions F_i and D_i. As long as both players play F_i the state remains s_1. Once either of them plays D_i the game moves to state s_2 and stays there. In this state each player has only one available action D_i. The payoff is depicted in Figure 2. For convenience we assume that after T periods the game does not end, but rather the state merely becomes s_2. The payoff is the sum of the stage payoffs.

	F_2	D_2
F_1	$1, 1$	$-2, 2$
D_1	$2, -2$	$0, 0$

$[s_1]$

	D_2
D_1	$0, 0$

$[s_2]$

Figure 2: stage payoff

This game can be described as a "slow" play of the TD: at the beginning it is agreed that the luggage is worth at least \$0, and at the first stage the players are asked: Is it worth more? If both say "yes," then it is now agreed that it is worth at least \$1, and at the next stage the players are asked: Is it worth more? This goes on for the prescribed number of stages T, or until somebody answers "no." Thus, the strategy of claiming k in the TD is equivalent to the strategy of saying "yes" (if you are asked) exactly k times, and thereafter say "no."[5]

Let σ be an equilibrium of the game Γ_T. Without loss of generality assume that σ is pure. For $1 \leq k \leq n$, let

$$A_k^i(\sigma) = \{\omega \in \Omega : a_k^i(\sigma, \omega) = D_i\}$$

Denote $P_k^i = 1 - \mathrm{P}(A_k^i(\sigma))$, and $P_{n+1}^i = E\left([T-n]_+\right)$.

For $0 \leq k \leq n$, let

$$g_k^i(\omega) = \sum_{t>k} h^i(a_t(\omega, \sigma))$$

i.e., it is i's payoff at the stages following k. For $1 \leq k \leq n$, let $\mathcal{B}_k^i$ be the σ-field generated by the following information: i's signal $\mathcal{B}^i$ at the start of the game, the play before stage k, and the state at stage k. Let

$$B_k^i = \{\omega \in \Omega : E\left(g_k^i \mid \mathcal{B}_k^i\right)(\omega) < 1\}$$

[5]From here on the proof is very similar to the proof of Theorem 2 in Neyman (1999).

and note that $B^i_k \subset A^i_k$ (otherwise, i is better off deviating from σ by playing D_i at stage k, when $\omega \in B^i_k \setminus A^i_k$). Hence, $\mathrm{P}(B^i_k) \le 1 - P^i_k$. Now, $E\left(g^i_k\right) = E\left(E\left(g^i_k \mid \mathcal{B}^i_k\right)\right) \ge \mathrm{P}[E\left(g^i_k \mid \mathcal{B}^i_k\right) \ge 1] = 1 - \mathrm{P}(B^i_k) \ge P^i_k$.

On the other hand, $h^i(a_t) \le 2$ for any a_t, and if $\omega \in A^{-i}_k(\sigma)$ then $h^i_k(a_t(\omega,\sigma)) \le 0$. It follows that for $0 \le k \le n$, $E\left(g^i_k\right) \le 2 \cdot \sum_{t=k+1}^{n+1} P^{-i}_t$. So we get

$$P^i_k \le 2 \cdot \sum_{t=k+1}^{n+1} P^{-i}_t$$

By induction on $(n-k)$, we get that for $0 \le k \le n$ and $i = 1,2$

$$\sum_{t=k+1}^{n+1} P^i_t \le 3^{n-k} E\left([T-n]_+\right)$$

Therefore,

$$E\left(g^i_0\right) \le 2\cdot\sum_{t=1}^{n+1} P^{-i}_t \le 2\cdot 3^n E\left([T-n]_+\right) \le 2\cdot 3^n \exp(-Kn) = 2\exp((\ln 3 - K)n)$$

By choosing $K > \ln 3$, we get that $\exp((\ln 3 - K)n) \to 0$ as $n \to \infty$. On the other hand $E\left(g^i_0\right) \ge 0$, because it is simple to guarantee a payoff of 0, and σ is an equilibrium. Therefore, $E\left(g^i_0\right) \to 0$ as $n \to \infty$. □

4 Remarks

Note that the above n is any number (not necessarily close to $E(T)$). So in fact we get

$$E\left(g^i_0\right) \le \inf_{n\ge 0}\ 2 \cdot 3^n E\left([T-n]_+\right)$$

In particular, if there exist $K > \ln 3$, and an infinite sequence $n_1 < n_2 < n_3 \ldots$ with $E\left([T-n_i]_+\right) \le \exp(-Kn_i)$, then $E\left(g^i_0\right) = 0$.

Neyman (1999) presents finitely repeated games, in which the number of repetitions T is not common knowledge. Among other things he shows that if the scale of uncertainty is smaller than some exponential function, then backward induction would still apply for the finitely repeated prisoner's dilemma. We take this opportunity to note that the proof given there actually proves a little more than stated. Similarly to the above paragraph, it is in fact necessary that the condition holds simultaneously for all integers, in order to avoid backward induction. That is, the proof in Neyman (1999) actually proves, in particular, the following compact sufficient condition for backward induction in the finitely repeated prisoner's dilemma.

Proposition 4.1. *If* $\inf_{n\geq 0}\left(2^{n+1}\cdot E\left([T-n]_+\right)\right)=0$*, then in equilibrium both players play Defect at every stage.*

In particular, if there exists a number $K>1$, and an infinite sequence $n_1<n_2<n_3\ldots$ with $E\left([T-n_i]_+\right)\leq 2^{-Kn_i}$, then the players defect at all stages.

4.1 Elasticity of the Traveler's Dilemma

We now compare the uncertainty about the upper bound T with uncertainty about the game's payoffs, which is used in the definition of the concept of elasticity in Bavly (2011). Specifically, we will translate the uncertainty about T to "elasticity uncertainties," which will enable us to conceive of the uncertainty about T as a subset, in a sense, of those uncertainties.

Consider an uncertainty structure about T, where each state ω is associated with an integer $T(\omega)$, the upper bound. This $T(\omega)$ induces the payoff function $u(\omega)(\cdot,\cdot)$, as we have seen. In this way the very structure represents uncertainty about payoffs, and we may think of the structure as such, and "forget" about $T(\omega)$. Or rather, we say that $T(\omega)$ is not a formal part of the description of a state of the world ω, but merely a concise way of describing the payoff $u(\omega)$, and the payoff is a part of a state's formal description.

Now the set of actions is still all integers, which is what was needed for talking about T. But as the game with this infinite set of actions is isomorphic to the game with actions from 0 to n (i.e., they have the same reduced form), Corollary 5.4 in Bavly (2011) enables us to move back to this original finite game.

Thus, we can deduce the exponential growth of the elasticity of the n-ranged traveler's dilemma from Corollary 2.2.

Theorem 4.2. *There exists a number $A > 1$, such that the elasticity of $(TD)_n$ is larger than A^n.*

We should note that we cannot use Theorem 2.3 to deduce that there is an exponential upper bound on the elasticity: the theorem tells us that such a bound holds for all structures with uncertainty about T; but the image of all these structures, after they are translated, is just a subset of all the payoff uncertainty structures.

By Proposition 6.1 in Bavly (2011), which enables translating back and forth between elasticity and "rationality elasticity," we also get this growth rate for rationality elasticity. We feel that this particular framework of perturbation from common knowledge of rationality is very appealing in the traveler's dilemma paradox. In other words, perhaps the analysis that leads to choosing $(0, 0)$ seems so unintuitive because we do not believe that the players really maintain 20 or 50 levels of mutual knowledge of rationality.

References

[1] Basu, K. (1994) The Traveler's Dilemma: Paradoxes of rationality in game theory, *American Economic Review* **84**, 391–395.

[2] Bavly, G. (2011) Elasticity of games, DP 592, Center for the Study of Rationality, Hebrew University, Jerusalem.

[3] Neyman, A. (1999) Cooperation in repeated games when the number of stages is not commonly known, *Econometrica* **67**, 45–64.