Some Thoughts on the Traveler's Dilemma

by Guillaume Alain

Abstract

In this paper, we will discuss the traveler's dilemma, a problem famous for its ability to point out how theoretical results of game theory don't always translate properly into real life applications. We explore the paradox of rational agents being outperformed by agents playing randomly and explain what is the problem with backward induction. We comment on the results obtained by [CGGH99] when carrying out the experiment with real subjects and argue that their modeling of the decision mechanism with randomness may lead to a reasonable fit with the experimental data, but it is of no real use. We point out how the game could be modified by adding randomness in a slightly different fashion to shift away the Nash equilibrium from its lose-lose position. Finally, we show how a simple mechanism of natural selection can provide insight into why people in real life don't play as "rational" agents and are influenced by the reward parameter.

1 Introduction

Game Theory features a good collection of toy example games in which the theoretical results are completely different from the way in which the games are played in real life. These games are sources of paradoxes when we try to reason about what the agents *should* do in a context where some of the parameters are unknown. Assuming that agents will go through a particular reasoning sometimes simplifies the game, but we can never know for certain what the others agents will do. They might play randomnly, and we would be wrong to assume that such a scenario is impossible in reality.

In 1994, Kaushik Basu introduced his own pathological game in [Bas94], called the traveler's dilemma. One of the special features of his game was that it was a "one-shot" game, as opposed to a repeated game like the centipede (see [Ros81]) in which an agent was given an opportunity to ajust his strategy when facing an opponent playing "irrationally".

In 1999, [CGGH99] presented their experimental results for the repeated game version of the traveler's dilemma, played by economics students with real money rewards. Not surprisingly, the game was played very differently than how theory predicted. The authors offer a statistical model to explain their data.

In this paper, we discuss the traveler's dilemma and explain the backward induction reasoning that leads to the disastrous Nash equilibrium. We have a look at the results of [CGGH99] and show some of the possible modifications that can be done to the traveler's dilemma in order to explain why agents would play anything else than the Nash equilibrium.

We explore in section 2 the possibility of adding randomness to the actions of the agents. Finally, in section 3 we will run a little simulation of evolution in which the agents will play the traveler's dilemma to determine which individuals are the fittest.

1.1 Statement of the problem

Two tourists fly home, each bringing back in his luggage a copy of an artifact bought as souvenir. The luggage are lost and the two tourists make claims to their insurance firm. The firm, wanting to discourage overvaluations of the lost artifacts, tells the two tourists that, if their two claims don't match, they will get the minimum of the two amounts and some reward R will be given to the tourist who provided that lowest of the two claims. That reward will be imposed as penalty to the other tourist.

The actual parameters chosen to define the game are not important, so before we introduce variations to the problem, we will stick to integer valuations in $\{2, \ldots, 100\}$ with a reward of $R = 2$ for any agent specifying a lower value and a penalty of $-R$ to the other.

1.2 The disastrous Nash equilibrium

The reasoning behind iterated removal is that if we know that no opponent will play a value $v \geq d \in \mathbb{N}$, playing $d-1$ dominates playing d, provided that $d-1$ is allowed to be played. Assuming that both players follow this reasoning to its conclusion, we get a backward induction that prunes every possible value except the smallest.

The flaw in that reasoning is that we assume that our opponent will actually reason this way. When the exposition can be simplified by putting ourselves in the shoes of an agent, we will be the first agent and our opponent will be agent 2 or player 2. If our opponent plays $v_2 = 100$, it's not true that, for us, playing $v_1 = 98$ dominates $v_1 = 99$ because it performs less well against that $v_2 = 100$. If there was no chance that our opponent would play $v_2 = 100$, then $v_1 = 99$ would be dominated, but otherwise it isn't.

For the purposes of this discussion, we will say that "rational agents" are agents who play the Nash equilibrium strategy and "irrational agents" are those who don't.

Suppose that we have a setting with two rational agents are playing $v = 2$ and one irrational agent is playing $v = 100$ all the time. If we do a round of traveler's dilemma with a rational agent, matching him with either of the others with equal probability, that agent has an expected utility of $\frac{1}{2}(2+4) = 3$. In the case of the irrational agent, his expected utility is $\frac{1}{2}(0+0) = 0$.

Suppose now that we have another irrational agent in the mix. Again we study random pairings. The rational agent has expected utility of $\frac{1}{3}(2+2+4) = \frac{8}{3}$, but the irrational agent has expected utility of $\frac{1}{3}(0+0+100) = \frac{33}{3}$, doing quite better.

Being crazy pays when you're not the only one.

Of course, one could argue that if the so-called rational agents in the two preceding contexts knew that they had an irrational overvaluing agent with them, they would play accordingly and move away from the Nash equilibrium $(2,2)$. We should also note that, if our opponent's action was a random variable of known probability distribution (equivalent to saying that we know his mixed strategy), we could easily compute our best response.

The point of the traveler's dilemma is that it's impossible to know beforehand what kind of "personality" our opponent will have. The richness of game theory has much to do with the fact that the behavior of agents is more complex than a uniform discrete distribution on the set of available actions.

1.3 The game played out in real life

In [CGGH99] is presented an actual traveler's dilemma experiment carried out with different rewards. They allow integer values ranging from 80 to 200 (paid in pennies) and they play repeated games of the traveler's dilemma with rewards of either $R = 10$ or $R = 80$ pennies.

Their conclusions support what we would normally expect from such a game (see figure 1. When rewards are low, subjects are more inclined to declare large values. When the reward/penalty is high (80 pennies), subjects cling to the Nash equilibrium. When facing an opponent playing 80 consistently, deviating from 80 to 81 leads to a variation in payoff from 80 pennies to none. For some fixed beliefs that we have concerning the probability distribution of our opponent's actions, our decision to play that Nash equilibrium or to deviate becomes purely a function of the reward.

When playing the game, it is tempting to do a few more steps of iterated removal in order to "outsmart" the opponent who might do less such steps. Doing more steps of iterated removals is not like multiplying ever-larger numbers mentally; once a subject figures out the trick of iterated removals, he is faced with an arbitrary decision of chosing how many he feels should suffice.

Moreover, it's important to note that the iterated removal reasoning is independant of R, yet as figure 1 clearly shows, this is not the case in practice. This rules out the possibility of modeling the

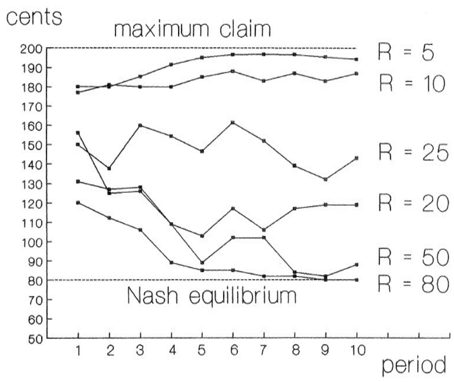

Figure 1: We include here a figure from [CGGH99]. They played repeated games of the traveler's dilemma changing partners randomly at every game. We study only one-shot games , but the graph is very informative about the relation between the declared values and the reward R.

personal preferences of people concerning the depth of their iterated removals outside a particular context.

Finally, we should note that the sums of money used in the experiment are not of the same magnitude as those of the original statement of the problem. We could certainly imagine how, if instead of pennies we were dealing with millions of dollars, agents could play very safely in order to secure 2 million dollars. The non-linear relation between money and utility is nothing new in economy, and we can only speculate about the way in which the results of [CGGH99] would differ if the sums involved in the game were scaled. This would fall more into the domain of psychology than game theory.

2 Randomness to break the disastrous Nash equilibrium

[CGGH99] offer a model for the repeated traveler's dilemma game where the agents are individually modeling the distribution for the claims of the other agents based on what they saw them playing. Agents then consider what are the expected utilities for their own actions in the next round, and make a decision with a rule that involves randomness but clearly favors the actions that have the highest utilities.

Their model is a logit function with a scaling parameter μ that determines how much the agents insist on favoring high expected utility actions. They do a fit for that μ based on their experimental data and find a value that appears somewhat reasonable by virtue of the resulting curve passing between the 6 points provided by their experimental data.

We believe that such a fit could have been done with many other statistical models and that, since their model itself does not appear to us as the "natural" choice to capture the essence of the cognitive process by which the subjects make their decisions, its usefulness is questionable. To validate their model they would have to make good predictions for another paradoxical game (possibly the centipede game, for example).

To illustrate our point, we present here a variation on the traveler's dilemma in which we are able to push the Nash equilibrium away from $(2, 2)$ by adding a certain degree of gaussian noise to the claims made by the agents. The original game used integers values, but we allow real values here.

One of the defining parameters of this new game is the variance σ^2 of the gaussian noise. The agents specify their actions $a_i \in [2, 200]$ and the corresponding outcomes are distributed as $o_i \sim \mathcal{N}(a_i, \sigma^2)$, rounding to the closest value inside $[2, 200]$.

We did a numerical simuation for various values of σ^2 to see what was the best response for player 1 in a context where there was no restriction on the claims, assuming that player 2's action is to claim 0 (we chose 0 for simplicity). When the best response is positive, it means that player 1 has incentive to play higher than player 2. When it's negative, player 1 has incentive to undercut player 2. This tells us on which side of the interval $[2, 200]$ the Nash equilibrium will be attracted. Note that, in some sense, the original noiseless traveler's dilemma corresponds to the case where $\sigma^2 = 0$.

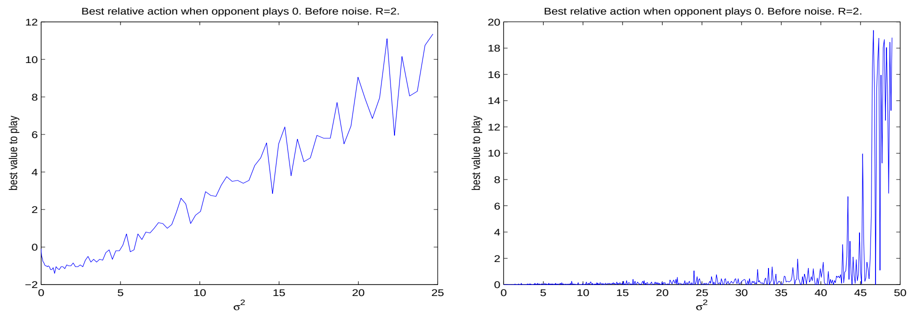

Figure 2: The left figure shows that, when the agents are playing values away from the boundaries of $[2, 200]$, as σ^2 grows they have incentive to play ever-rising values. On the right figure, the same numerical analysis is done by rounding to 0 all the negative values resulting from the gaussian noise. We can see that for σ^2 sufficiently large the agents at the lower possible value too have incentive to deviate upward.

As expected, we can see that when σ^2 is very small, the best response is to play very closely under 0. With σ^2 sufficiently large, however, the best response is now to play higher than what the other agent is playing. This is consistent with the fact that any agent wishing to play low values has to accept that he will often miss opportunities to benefit from an opponent playing high values "accidentally" (because of the noise). This curve should be smooth, but the roughness is due to randomness in the numerical simulation.

For those large values of σ^2 for which the agents have incentive to play large numbers, we don't really know the exact form of the new Nash equilibrium, but we think that it involves a good amount of weight on high values. We didn't study this further as our goal was qualitative and not quantitative.

The point of this section is simply to argue that the Nash equilibrium can be lifted from $(2, 2)$ by adding some randomness in the game. When rational agents are playing the game, they have to be Bayesian about the value of σ^2. If the σ^2 was known in advance and it was large enough, every agent would play a high value.

This is nice, but when the agents know that no randomness will be added to their actions, why would they reason like this and assume that σ^2 could be large ?

3 Numerical simulation of evolution

In this section we play with a simple model of evolution to argue that in the context of natural selection, agents have to take risks in order to maximize their chances of survival. When the lower 20% of the population gets weeded out at every game and replaced with a slightly mutated version of the top 20%, playing the low Nash equilibrium strategy is a good way to end in the lower 20%. Undercutting the other agents by 1 is still a good thing to do, though.

We return to the setting in which the allowed values are integers in $[2, 200]$. Agents always play the same value. They never adapt their strategy, but when agent i is chosen to generate a replacement for one of the killed agent j, the new value for agent j is picked from $\mathcal{N}(a_i, \sigma^2)$ and rounded to the closest integer in $[2, 200]$.

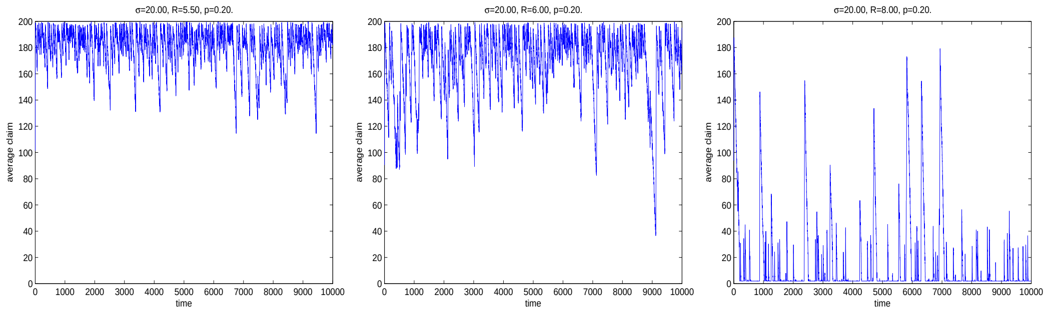

Figure 3: Average claims of the fittest agents over time. We look at $\sigma = 20, p = 0.20$ and $R = \{5.5, 6.0, 8.0\}$ in that order. We can see that when the reward for undercutting increases sufficiently, the agents are encouraged to play lower values.

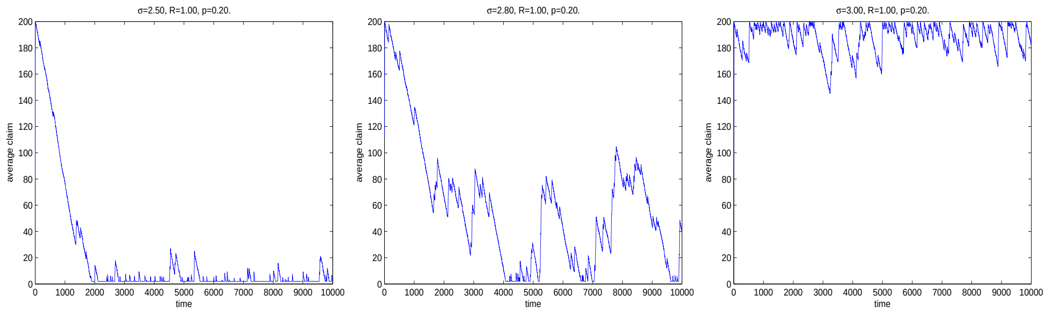

Figure 4: Average claims of the fittest agents over time. We look at $R = 1, p = 0.20$ and $\sigma \in \{2.5, 2.8, 3.0\}$ in that order. We can see that when the variance is sufficiently large, the agents are encouraged to play high values.

The parameters for the simulation are

$$\begin{aligned} n &: \text{ the number of agents in the population} \\ p &: \text{ the proportion of agents killed and replaced at every iteration} \\ \sigma^2 &: \text{ the variance parameter for the mutations} \\ R &: \text{ the reward in the traveler's dilemma.} \end{aligned}$$

We are interested in the average claims made over time by the surviving agents. There are essentially two different kinds of behaviors to look for : when the surviving agents play the lowest possible values and when they don't. The latter is more interesting than the former. We present in figures 3, 4 some graphs for particular choices of parameters that illustrate the fact that, when there is sufficient mutation variance (compared to the reward), it drives the population to high values. The level of selective pressure (increasing as a function of p) can also be studied (see figure 5). Tweaking both the variance σ^2 and the reward is somewhat redundant because we only want to observe if, when one of the values is fixed, there is a threshold value for the other such that the behavior of the agents changes qualitatively. We still included the two flavors here.

An interesting feature of the cases in which the agents play high values is that there are cycles involved. The fittest agents are those undercutting the other agents by a hair (remember that agents never alter their values), so we can expect the average claim to go down over time. However, the little touch of randomness in the mutations makes it such that it sometimes becomes viable to play higher values and the conservative agents who undercut are eliminated.

The graphs suggest some form of convergence, but we believe that they actually represent only general tendencies of distributions, and that even in the best scenarios, there will always be a (very unprobable) spiral to the bottom followed by a (highly probable) raise back to the high values.

A valid objection could be that the use of randomness in the mutations is no different from the use

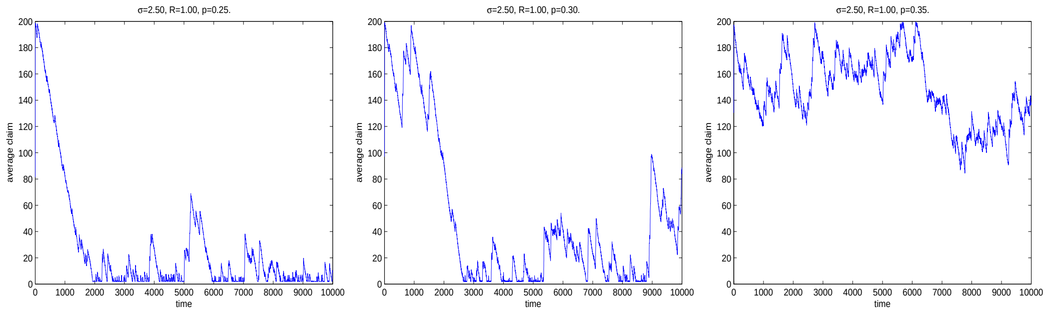

Figure 5: Average claims of the fittest agents over time. We look here at the effets of the level p of natural selection. We use $\sigma = 2.5, R = 1$ and $p \in \{0.25, 0.30, 0.35\}$ in that order. We can see that when more agents are killed, the incentive to take risks to stand out increases.

of randomness in the noisy actions scenario from section 2. It makes a difference from the point of view of the agents, however, because they need not be aware of the actual variance used in the reproduction mechanism in order to play optimally in terms of survival. They can assume that Nature tuned their definiting parameters $\{a_i\}$ according to what was the most efficient given the exact value of the variance σ^2 that only Nature knows. Following their instincts is the best course of action.

4 Conclusion

The cold rational analysis of the traveler's dilemma is an interesting exercise, but it's of little value in real life settings because it fails to consider the possible presence of so-called "irrational" agents.

However, accounting for that would require us to model complex human emotions like the sense of fairness, justice, the value of money, the sense of sharing, all of which are very dependant on the setting in which the game is played. As show by [CGGH99], the reward parameter influences much the agents. The results would vary even more if we were to conduct such an experiment in different regions of the world at different times. We think that the best that can be done is limited to a posterior curve fit with a model that can "explain" any set of experimental results, but not predict new results accurately in a different game setting, even though the same parameters of human irrationality are involved.

References

[Bas94] Kaushik Basu. The traveler's dilemma: Paradoxes of rationality in game theory. *American Economic Review*, 84(2):391–95, May 1994.

[Cam03] C. Camerer. *Behavioral game theory: experiments in strategic interaction.* Princeton University Press Princeton:, 2003.

[CGGH99] C.M. Capra, J.K. Goeree, R. Gomez, and C.A. Holt. Anomalous Behavior in a Traveler's Dilemma? *The American Economic Review*, 89(3):678–690, 1999.

[Ros81] R.W. Rosenthal. Games of Perfect Information, Predatory Pricing and the Chain-Store Paradox. *Journal of Economic Theory*, 25(1):92–100, 1981.