Analyzing
The Traveler's Dilemma

Multi-Agent Systems Project

Sander Land Joël van Neerbos Tim Havinga

April 10, 2008

Chapter 1

Introduction

The traveler's dilemma [1, 2], is described with the following parable.

Two travelers returning home from a remote island, where they bought identical antiques (or, rather, what the local tribal chief, while choking on suppressed laughter, described as "antiques"), discover that the airline has managed to smash these, as airlines generally do. The airline manager who is described by his juniors as a "corporate whiz", by which they mean a "man of low cunning," assures the passengers of adequate compensation. But since he does not know the cost of the antique, he offers the following scheme.
Each of the two travelers has to write down on a piece of paper the cost of the antique. This can be any value between 2 *units of money and* 100 *units. Denote the number chosen by traveler i by n_i. If both write the same number, that is, $n_i = n_j$, then it is reasonable to assume that they are telling the truth (so argues the manager) and so each of these travelers will be paid n_i (or n_j) units of money.*
If traveler i writes a larger number than the other (i.e., $n_i > n_j$), then it is reasonable to assume (so it seems to the manager) that j is being honest and i is lying. In that case the manager will treat the lower number, that is, n_j, as the real cost and will pay traveler i the sum of $n_j - 2$ and pay j the sum of $n_j + 2$. traveler i is paid 2 *units less as penalty for lying and j is paid* 2 *units more as reward for being so honest in relation to the other traveler.*
Given that each traveler or player wants to maximize his payoff (or compensation) what outcome should one expect to see in the above game? In other words, which pair of strategies, (n_i, n_j), will be chosen by the players?

In this paper, we will analyze the Traveler's Dilemma from several angles, and try some variations on it, like changing the reward. From these analyses we will try to explain this apparent contradiction between classical game theory and human behaviour.

Chapter 2

Game Theory and the Traveler's Dilemma

The traveler's dilemma was originally created as an example of a conflict between intuition and game-theoretic reasoning. In this chapter we will explore what game theory can be used to analyze the game, what basic assumptions it makes about the players' behaviour, and how this relates to actual human behaviour.

In this chapter we will use $\pi(p, q)$ to denote the payoff for player one (using strategy p) playing against player two (using strategy q). For the traveler's dilemma, this function is:

$$\pi(p,q) = \begin{cases} p & p = q \\ p+2 & p < q \\ q-2 & p > q \end{cases}$$

2.1 Nash Equilibria

One very important tool in game theory is a so-called Nash equilibrium [3], it defines a kind of optimal strategy. It is informally defined as a set of strategies (one for each player), such that no player can do better by choosing another strategy while keeping the others' strategies fixed. That is, if all choices of all players are known to everyone, no one would want to be the only one to switch strategies. In a two player game, it would be a pair of strategies p, q such that $\pi(p\prime, q) \leq \pi(p, q) \quad \forall p\prime \neq p$ and $\pi(p, q\prime) \leq \pi(p, q) \quad \forall q\prime \neq q$.

A strict Nash equilibrium is a set of strategies such that any player will do strictly worse by switching, so in a two-player game it would be a pair of strategies p, q such that $\pi(p\prime, q) < \pi(p, q) \quad \forall p\prime \neq p$, etc.

This situation is not always (in fact, usually isn't) the optimal outcome for everyone. For example, in the prisoner's dilemma, the (strict) Nash equilibrium would be both players defecting, but the optimal outcome is both players cooperating. However, cooperation is not a Nash equilibrium because unilaterally defecting is a locally better strategy.

Likewise, in the traveler's dilemma a Nash equilibrium is $(2, 2)$. No player can do better by choosing a higher value, because they would be picking the higher value, and receive \$0 as a result of the other player's choice of 2. This is also the only Nash equilibrium in the game, because with any other pair of strategies $(x, y) \neq (2, 2)$, at least one player can improve their strategy by picking a value that is exactly one lower than the other player.

2.1.1 Rationalizable Nash Equilibria

There is also the concept of a rationalizable Nash equilibrium. This involves repeated elimination of *dominated strategies.* A (weakly) dominated strategy is a strategy p for which another strategy q exists such that $\forall x\ \pi(q, x) \geq \pi(p, x)$ (and at least one x for which strict $>$ holds). That is, q is a strategy with at least an equal payoff, regardless of the other player's strategy, and is strictly better for at least one case. To find the rationalizable Nash equilibria, all dominated strategies are eliminated for both players (equivalent to an assumption of some rationality in both players), and this process is then repeated with the remaining strategies until no changes occur. All remaining strategies are then rationalizable Nash equilibria. The assumption of rationality made here seems like a very reasonable one, when given the choice between 'win \$100' and 'win \$100 for sure, and possibily \$101', almost anyone would choose the latter. However, the iterative process doesn't seem like something the average player would intuitively apply in his reasoning, so this may cause problems, as we will see later on.

Also note that in one experiment with game theory experts [9], one fifth of all players played the strategy \$100, which is strictly dominated even without eliminating any strategies, so maybe even this very basic assumption of rationiality is too big of an assumption even without the iterative elimination of strategies.

In the traveler's dilemma, $(2, 2)$ is also the only rationalizable Nash equilibrium, because, as we can see in te following table, the choice of \$100 is strictly dominated by \$99:

x	$\pi(100, x)$	$\pi(99, x)$
100	100	101
99	97	99
$n < 99$	$n - 2$	$n - 2$

Eliminating the strategy of playing \$100 for both players, in the remaining choices of \$2 – \$99, the choice of \$99 is strictly dominated by \$98, according to similar reasoning. This process of elimination can be repeated until the choices are \$2 – \$3, where the choice of \$3 is strictly dominated by \$2, leaving \$2 as the only rational choice for both players.

Finding this rationalizable Nash equilibrium also clearly shows the reasoning and assumptions necessary for a player to get to the conclusion that playing \$2 is the 'best' strategy.

- I am a rational player, so I will not play the dominated strategy \$100.
- You know I am a rational player, so you will not play the dominated strategy \$99.
- I know you know I am a rational player, so I will not play the dominated strategy \$98.
- You know I know you know I am a rational player, so you will not play the dominated strategy \$97.
- and so on until...
- (I know you know) $\times$ 48 I am a rational player, so I will not play the dominated strategy \$4.
- You know (I know you know) $\times$ 48 I am a rational player, so you will not play the dominated strategy \$3.
- I know this, therefore I will play \$2.

The assumption is also clear, it has to be *common knowledge* that a player will not play a dominated strategy. This seems like a very reasonable assumption at first, but when looking at the above steps it is clear that hardly anyone would reason that many levels deep. Indeed, it is practically impossible to imagine the difference between "you know (I know you know) $\times$ 10" and "you don't know (I know you know) $\times$ 10", let alone repeating this for another 80 steps.

Kripke models

If we would try to model the situation in a Kripke model, we would need a lot of worlds. Because both players initially have no real knowledge about what choice the other player will make, they should consider every legal number (every number from 2 to 100 in our case) as a possible choice of the other player. This has as a consequence that the needed amount of worlds to model the situation would be the square of the number of legal choices (99 in our case, with which we would get a total number of 9801 worlds). A Kripke model with this much worlds would not serve much of a purpose. To reduce the amount of worlds in the Kripke model, we greatly reduced the number of legal number choices:

every player can now only choose between the numbers 2, 3 and 4. This gives the Kripke model shown in Figure 2.1.

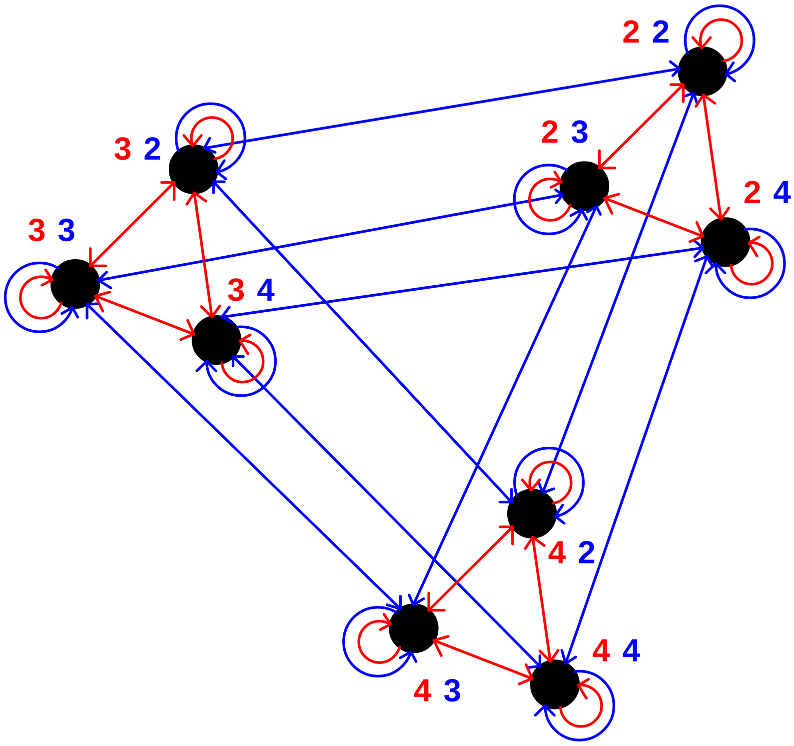

Figure 2.1: The Kripke model of the situation where both players can choose between the numbers 2, 3 and 4. The choices and relations of player 1 are shown in red, the choices and relations of player 2 are shown in blue.

In this figure we can see that both players have in each world arrows to all worlds where the choice of the player itself is the same, but where the choice of the other player is different. This is because a player obviously knows which number it chooses, but does not know which number the other player chooses.

If, in this situation, it would become common knowledge that choosing 4 is not rational because it means playing the dominated strategy \$4, and it is common knowledge that both players are rational, then the Kripke model can be reduced to the Kripke model shown in Figure 2.2.

As can be seen by comparing this model to the previous one, this model is the previous model but without the worlds where one of the players would choose 4.

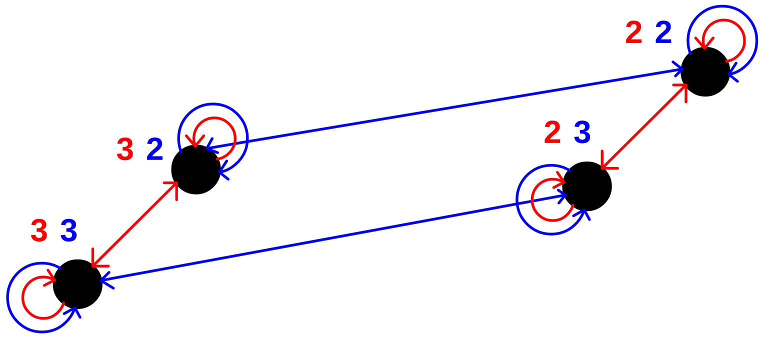

Figure 2.2: The Kripke model with only the choice between 2 and 3. The choices and relations of player 1 are shown in red, the choices and relations of player 2 are shown in blue.

Similar to removing all world with choice 4, all worlds with choice 3 can also be removed, as playing $3 has now become a dominated strategy. This gives the Kripke model shown in Figure 2.3.

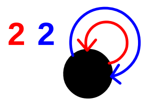

Figure 2.3: The final Kripke model, with no choices at all. Both players will choose the number 2, and this is common knowledge.

In this Kripke model, it is common knowledge that both players will play $2. This is caused by the assumption that both players are rational and won't play dominated strategies, as described before. So in fact there is only one possible world for both players, because they know both their choice and the choice of the other player.

Another way to model the dilemma could have been to just have one world where both players are rational and both players have reflexive arrows, so it is common knowledge that they are rational. This would not be an interesting Kripke model, but the model would be able to graphically support the "I know that you know that I know ..."-reasoning as described before.

Either way, it does not seem to be very useful to model the (entire) situation in a Kripke model. You would either have an enormous amount of worlds, or you could reason to the point that both players will play $2 and have only one world. Both cases do not seem to give more insight into the situation.

2.1.2 A bigger reward?

Increasing the honesty reward b given to the player with the lowest choice clearly does not affect any of the results. The (rationalizable) Nash equilibria are still at the minimum choice (to avoid negative payoffs, we will make this b, so the possible choices are $\$b$ – \$100). Even the reasoning needed to get to the Nash equilibrium stays the same, for example with a reward of \$10. In this case, the strategies \$91 – \$99 now all dominate \$100, the strategy \$99 is not dominated before \$100 is eliminated, because no other strategy has a possible payoff of $\geq$ \$109.

However, it is easier to be rewarded for an uncooperative strategy. When the reward is two, choosing a value more than two below the other player's choice is worse than choosing the same strategy as him, and being more than three below it even means the 'completely irrational' choice of \$100 would have been better.

Capra et al. [10] also looked at increasing the reward. They show that, as the reward increases, the players are more and more intended to go to the Nash equilibrium. For the low reward fees (from \$1), the average claim is close to the maximum claim. But as the reward increases, the players are willing to take a bit more risk for a higher payoff. And finally, for the largest values of the reward (with a maximum of \$40), the average reaches the Nash equilibrium.

They conclude that, however the game should not be influenced by the reward, according to the Nash equilibrium, in fact it is when tested in real life. This does not mean that the Nash equilibrium is a wrong method, only that it fails to explain human behaviour for this problem.

2.2 Evolutionary stable strategies

Closely related to the theory of Nash equilibria is the concept of an *evolutionary stable strategy* [3]. A strategy p is an Evolutionary Stable Strategy (ESS) if a pure population of individuals using strategy p is stable against invasion by any other strategy. That is, if a single $q-$strategists appears in a population of $p-$strategists, he will have a lower fitness than everyone else, so he will be outcompeted by the $p-$strategists, eventually resulting once again in a population consisting of only $p-$strategists.

Whereas in a Nash equilibrium no player can benefit from switching, in a non-strict equilibrium a different strategy that has an identical payoff is still allowed. This is justified by the assumption that a player has no incentive to switch from the equilibrium, because it would not benefit him. However, in an evolving population there does not need to be an incentive (i.e. a fitness advantage) to switch, because genetic drift is sufficient to cause divergence from the equilibrium in the absence of a penalty for doing so.

Therefore, a stricter concept is needed, one of these is the evolutionary stable strategy. Originally developed by John Maynard Smith [4] for use in theoretical biology , it has also been used in game theory, economics and many other fields of research.
It is defined as a strategy p for which for all other strategies q the following conditions hold:

- $\pi(p,p) > \pi(q,p)$ *or*
- $\pi(p,p) = \pi(q,p)$ and $\pi(q,q) < \pi(p,q)$

The first condition is the same as for the strict Nash equilibrium, so all strict Nash equilibria are evolutionary stable strategies. Maynard Smith's second condition signifies that even though the strategy q is neutral when playing against p, the population of players who continue to play strategy p still have an advantage over those playing strategy q, when meeting a q-strategist.

This definition shows that every ESS is also a (possibly non-strict) Nash equilibrium, but not every non-strict Nash equilibria is an ESS. Consider a game with the following payoff matrix:

	C	D
C	(1,1)	(0,0)
D	(0,0)	(0,0)

In this game (D,D) is a non-strict Nash equilibrium, but it is not an ESS, because any number of $C-$strategists could invade the population, and would then be at an advantage playing against each other.

	C	D
C	(1,1)	(1,1)
D	(1,1)	(0,0)

In this game (C,C) is a non-strict Nash equilibrium, and is also an ESS, because the 2nd condition holds: no more than one $D-$strategist could invade the population without experiencing a penalty for meeting other $D-$strategists.

For the traveler's dilemma, $(2,2)$ is also an ESS, simply because it is a strict Nash equilibrium.

Sometimes a stronger definition is used, proposed by Thomas [5], which stresses the importance of the payoff of invaders more: :

- $\pi(p,p) \geq \pi(q,p)$ *and*
- $\pi(q,q) < \pi(p,q)$

In this case, not all strict Nash equilibria are ESS. Also, any ESS according to this definition is also an ESS according to Maynard Smith's definition.

For the traveler's dilemma, $(2, 2)$ is an evolutionary stable strategy according to the first definition but not according to the second, after all $\pi(100, 100) > \pi(2, 100)$.

2.3 Simulation

Because we can not simulate human players, we will try to simulate a population of 'travelers' evolving, and see if the Nash equilibrium appears as a result.

2.3.1 Methods

We simulated a population of travelers using a co-evolutionary genetic algorithm [6]. A co-evolutionary GA does not use a fitness function to rate the various solutions, but rates their performance when several members of the population 'fight' eachother in some way.

We used the genetic algorithms library 'charlie' [7], for the Ruby programming language. Source code can be seen in the appendix. For the selection algorithm, we used the built-in `CoTournamentSelection`. In this algorithm, groups (tournaments) of size T were created, and all travelers play a single-shot traveler's dilemma against all others in the group. After all these games have been played, the group was replaced by a normal round of roulette selection, taking their total score over all these games as their fitness values. This means the fitness of an individual is its average performance in playing single-shot traveler's dilemma games when facing other individuals in the tournament group. In each generation, enough of these tournaments are played such that roughly the entire population is replaced once, so the generations listed here are technically 'generation equivalents'.

No crossover was used, since the genotype is only a single number. Mutation rate was set at $p = 0.5$, that is, half of the children are different from both their parents, and the other half are identical to one of their parents. Mutations were applied by adding a Gaussian distributed random number ($\sim N(0, s)$) to the gene, and then clamping it to the range $(\text{reward}, 100)$. The population size was set at the default 20 with random initialization, and the algorithm was run for 1000 generations.

The honesty reward was varied between 2 and 10, and for each different reward the GA was run for a variety of mutation and tournament sizes. For the tournament sizes we used the full possible range $T = 2, 3, \ldots, 19, 20$, and for the mutation sizes we used $s = 0.1, 0.2, \ldots, 9.9, 10$, so in total 1900 runs were done for each different reward.

We repeated all these tests using a population which was initialized at the minimum value possible, which is the honesty reward. This way we could study the non-ESS results in the other tests to see where they were caused by a failure to converge within the limited number of generations, and where the population actually escaped from the ESS.

2.3.2 Results

Because of the vast amount of results (about 40 million generations) we can not present, or even examine, all the data. Instead we only look at the average move in the population at the end of a run, and visualized these results using MATLAB.

The images can be interpreted as follows: On the horizontal axis is the mutation size s, increasing from left to right and on the vertical axis is the tournament size T, increasing from top to bottom. Each single rectangle, of which there are $19 \times 100 = 1900$ in each image, represents the average strategy in the population after 1000 generations and is represented by a color which can be looked up in the bar to the right of the image. In general bright red represents a value close to the maximum 100, and blue a value close to the minimum value.

Figure 2.4 shows the results for the experiments with random initialization.

When using a tournament size of 2 the population converges to the minimum value for all settings. This can be easily explained by the fact that the fitness value in this case is just $\pi(p, q)$, i.e. those individuals who simply have a higher payout than their opponent, regardless of what this payout is, are more successful. For bigger tournaments, an individual can 'win' every game and still have the lowest average, for example a \$2 player in a population of \$100 players. These results then do not really represent what we want to model, which is maximizing your payout itself and not your payout relative to a single opponent. However, they can still be used as a baseline for the amount of convergence and amount of drift when there is a clear selective pressure for the lower strategies. Likewise, the bottom row, for a tournament with the entire population, is closest to the actual traveler's dilemma and therefore contains the most important results.

Looking at figure 2.4(a), it is clear that most runs other than the $T = 2$ results converge to a high value, somewhere in the range \$85 – \$100. For higher rewards, a front appears from which on one side (small tournament size and especially, low mutation) the results converge to the minimum value, and on the other side converge to a value near the maximum.

For larger reward, there are no longer any settings that don't converge to the minimum, although figure 2.6 shows this is just an effect of the maximum mutation size.

These experiments were repeated starting with populations with every individual at the minimum value, so right in the ESS. Maynard Smith's definition of an ESS suggests that it should be very unlikely for the population to escape the ESS, but as can be seen in figure 2.5, the results are very similar. Only the parts which had somewhat inconclusive results in the previous experiments now clearly stay at the ESS, but this did not because the boundaries shift much,

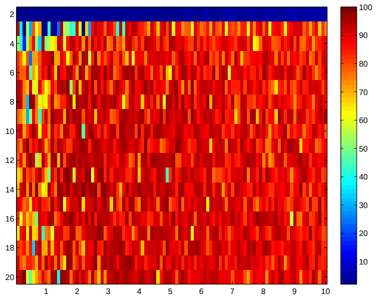

(a) $b = 2$

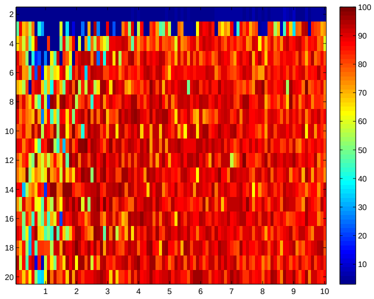

(b) $b = 3$

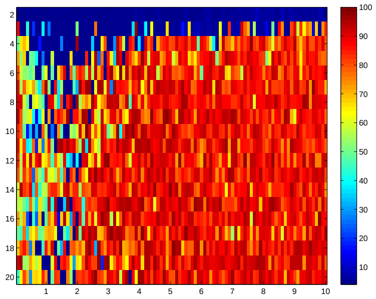

(c) $b = 4$

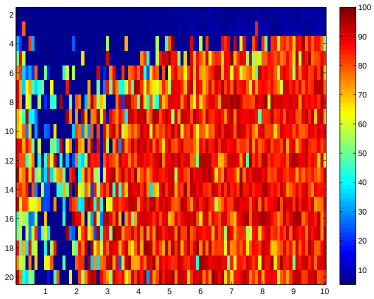

(d) $b = 5$

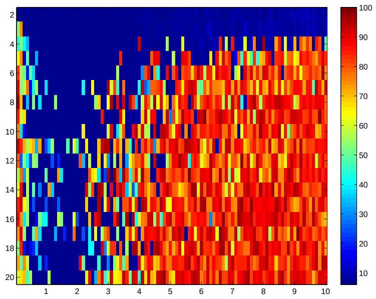

(e) $b = 6$

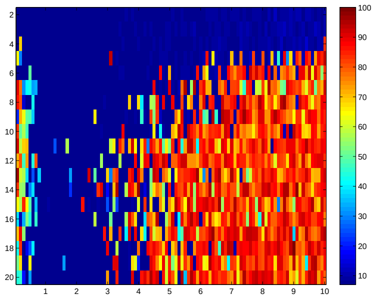

(f) $b = 7$

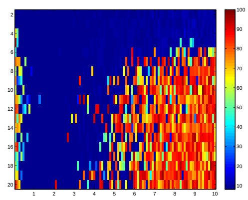

(g) $b = 8$

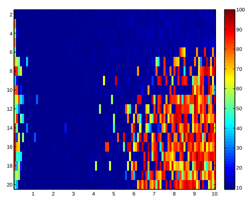

(h) $b = 9$

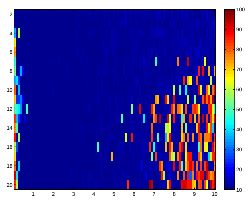

(i) $b = 10$

Figure 2.4: Results of the experiments with random initialization. The $x-$axis has the mutation rate and the $y-$axis shows the tournament size.

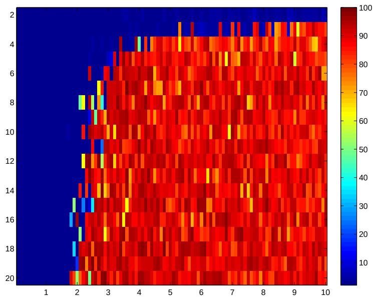

(a) $b = 2$

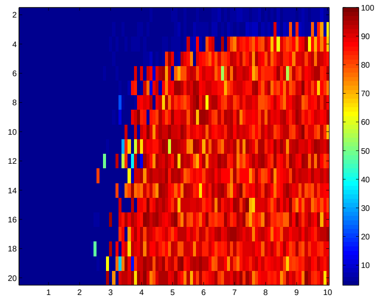

(b) $b = 3$

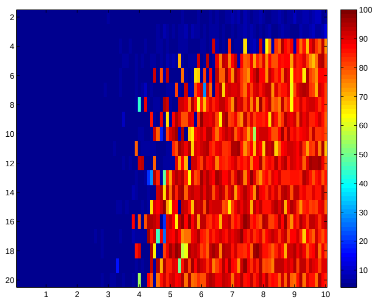

(c) $b = 4$

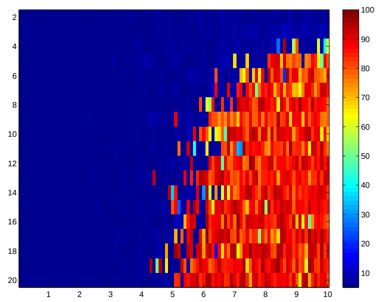

(d) $b = 5$

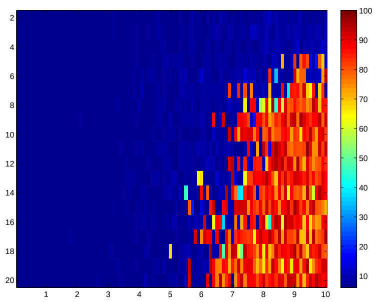

(e) $b = 6$

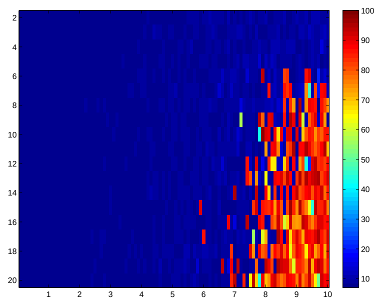

(f) $b = 7$

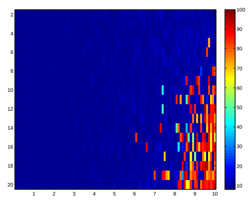

(g) $b = 8$

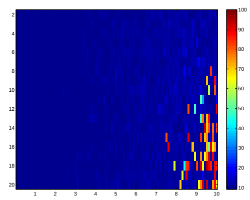

(h) $b = 9$

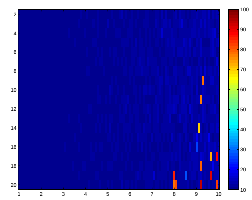

(i) $b = 10$

Figure 2.5: Results of the experiments with initialization at the ESS.

instead they just become sharper. For example, in the $b = 4$ figures in the bottom $T = 20$ rows, the boundary to the right of which all results are ≥ 75 is at about $3.7 - 3.9$, whereas in the ESS results it is at about 4.5.

Clearly, the second definition of an evolutionary stable strategy is the better one here, and also explains why the first definition fails.

Simple calculations show that when there is sufficient variation in the population, the strategy with the highest payoff is in fact above the average of the other strategies. This explains the selective pressure towards higher strategies, and with it possibly the strategies seen in real life.

2.4 Conclusion

Classic game theory fails to explain human behaviour when playing the traveler's dilemma, but can still be used to see what assumptions are being made. Common knowledge of rationality rarely exists, because most people are just not inclined to think more than a few levels deep. Even game theory experts don't assume this when playing against each other.

Instead of talking about rational or irrational choices, it may be better to

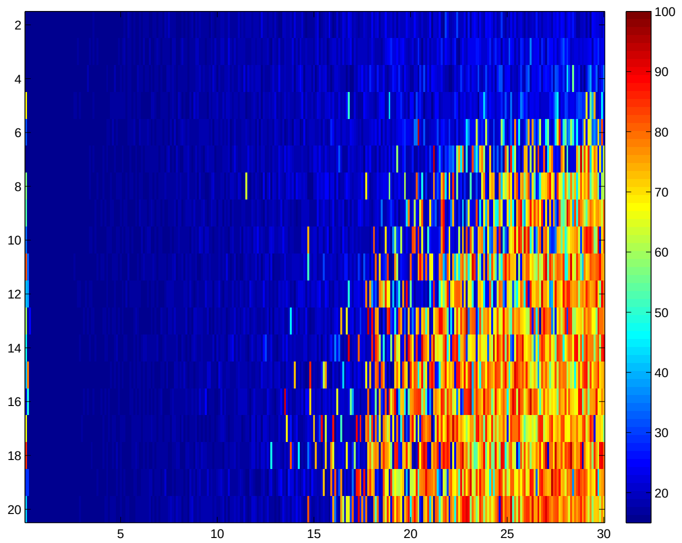

Figure 2.6: Results for a higher reward $b = 15$ with mutation sizes up to 30 and random initialization.

see the game as one with various levels of cooperation. We think most people would expect others to be cooperative (i.e. conspire against the insurance company), and would therefore pick a high value. Unlike some other games, even uncooperative players who know for sure their opponent is cooperative are best off picking a value above the average strategy of their opponents if they can not estimate these strategies to within a very narrow range. This basically forces them to be cooperative as the only optimal strategy.

Only when the reward for being uncooperative is high enough, and the variation in your opponents' strategies is small enough, is it possible to play an optimal strategy at the expense of all others, and will the strategies likely trend towards the Nash equilibrium.

Chapter 3

Real-life strategies

3.1 Theory of Mind

How humans may react can also be predicted using the paper by Verbrugge and Mol [8]. This paper suggests that people may not use high order Theory of Mind: they do not reason to the degree that their number of choice approaches 2. Or, they may assume that the other player does not use high order Theory of Mind, so the other player will not reason to that same degree. Of course it is also possible that both cases are true.

The paper also suggests that how pragmatic the player will be is dependent on how cooperative the player is. If the player wants to maximize the common profit, it will naturally choose a high number (probably 100). But if the player is uncooperative, for example if he or she has a strong desire to 'win' the honesty reward, then he or she will probably show less pragmatic behaviour, assuming having a higher claim is less useful than being the one that obtains the honesty reward.

Another situation that may occur is that one or both of the players will assume cooperative or even altruistic behaviour (to some degree) from the other player. This makes it more attractive to choose a higher number, even if the player has a desire to obtain the honesty reward.

3.2 Game Theory Experts

Another approach to see what the logical result of the Traveler's Dilemma would be, when played by human players, is given by Becker et al. [9]

They wanted the participants that play the game to make the most logical choice. Therefore, they addressed members of the illustre Game Theory Society. These would, of all people, make the most logical decision. To encourage this, one of the participants would receive a prize, equivalent with the performance of their strategy. And another, for the goodness of fit of their belief of the other participants' strategies.

Participation to the experiment consisted of two steps. First of all, all participants had to specify what their belief was about what the other participants were going to enter. This was done by giving a probability value to each value in the strategy space $\{2, 3, ..., 100\}$. Second, all participants had to specify their own strategy.

The participants had the choice to do two strategies: they could either play on a single value (a pure strategy), or assign certain probabilities to a number of values (a mixed strategy).

3.2.1 Results

51 participants entered a strategy, from whom 47 entered a belief of the average strategy.

The average choice of the participants in Becker et al.'s experiment is shown in Figure 3.1. The large averages at 31, 49 and 70 come from pure strategies with those values. As can be seen, some of the participants actually played the Nash equilibrium of 2. Also, 10 participants entered the 'irrational' strategy 100. The authors call this strategy irrational, because the strategy 100 is strictly dominated.

The Nash equilibrium was the least profitable strategy, as can be seen in Figure 3.2. Note here that the payoffs shown are directly dependant on the strategies of the other players. Therefore, the ultimate strategy differs per experiment.

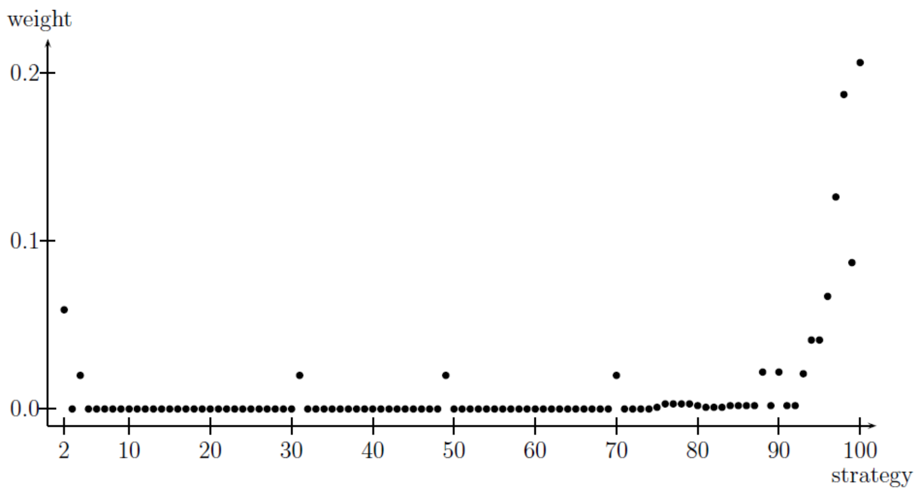

Figure 3.1: The average strategy, played by the participants in the experiment. Image taken from [9].

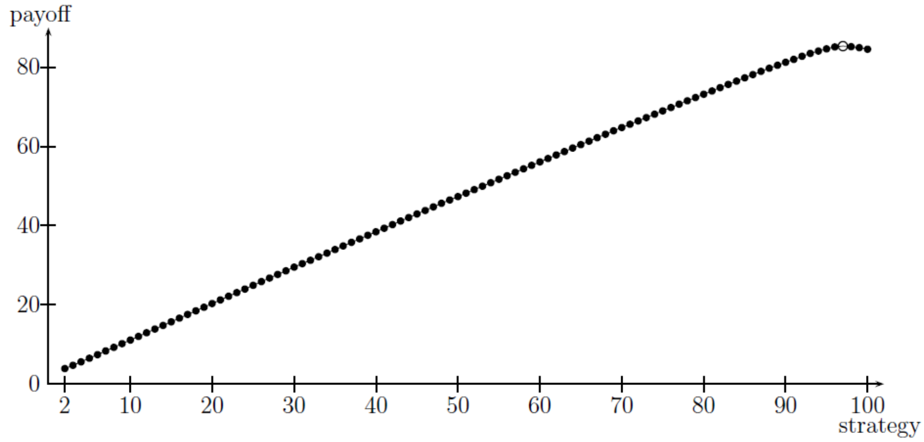

Figure 3.2: The payoff per strategy. As can be seen, the Nash equilibrium has the lowest payoff, and the pure strategy 97 has the highest payoff. Image taken from [9].

A nice result in the beliefs given by the participants, was that only 17 (out of the 47) of the participants played a strategy that coincided with their belief of the average strategy. From this we can deduce that some of the participants deliberately avoided playing the average strategy, to gain a higher payoff.

Now, Becker et al. claim that the Traveler's Dilemma becomes a game of incomplete information. This must be the case, because there is no single winning strategy. The only logical strategy that can be inferred is the Nash equilibrium, and they claimed that this is an irrational choice.

Next, Becker et al. disregard all strategies but those that fall (partially) into the interval $[94, ..., 99]$. Disregarding the players that chose the strategy 2 or the strategy 100, of the remaining 38 participants, 27 (more or less) adhere to this interval model.

After neglecting all strategies outside this interval model, a weight of 0.281 is not accounted for, the rest is contained in the model. They now think that the resulting model is a reasonably good approximation to the weight of each of these points becoming an equilibrium.

3.2.2 Conclusions

Becker et al. conclude that their resulting model is a good one for individual predictions, because of the fact that all participants were experts in game theory. They therefore conclude that the fact that their interval model is reasonably high up the strategy space, indicates that these high values are not because the participants did not understand the game (and therefore did not end up with the Nash equilibrium 2), but the high result is the result of a robust pattern of behaviour.

While analyzing the game as a game of incomplete information, the authors still have no arguments for the fact that almost 20 percent of the participants chose the, to them, irrational strategy 100. However, they note that, when no player would play the irrational strategy 100, all players would end up in the Nash equilibrium of 2.

We can reason that, when there are participants that play the strategy 100, there is no need for the other participants to go much lower than the honesty reward below that. As shown in the second chapter, the choice for 100 is strictly dominated, so there is no reason to take it. But if it is taken (for example as a cooperative strategy) then it would be foolish to go much lower, and get a lower payoff than you could have had.

Whilst if the strategy 100 was abandoned, the race for the just-a-bit-lower claim would start, and finally, all players would indeed end up in the Nash equilibrium of 2. But this reasoning is not mentioned by the authors.

We think that the way that Becker et al. handled the problem is curious. The way they analyze the data is good, but how they handle the observed data, and how they think about the contestants is remarkable.

First of all, they say some of the result is irrational, while they tried their best to let the participants be as rational as possible. Next, they discard part of the results, taking a range that fits them, and then saying that the results can be fitted well inside that range. While a more statistical approach would be to first try to analyze what the probable result would be, and then actually compare the results with the model created at first.

Chapter 4

Conclusion

In this paper we analyzed the Traveler's Dilemma. While analyzing the dilemma using classical game theory, we found the rationalizable Nash equilibrium in which both players choose the number 2. However, human players do not make this choice very often in practice, and even simulations done by computers seem to indicate otherwise.
Changing the reward for being honest (or for just being lucky to have chosen the lowest number) to a number higher than 2 should theoretically have no impact on the strategy a player will use, but in practice it turns out that it actually does have impact on the choices human players make. When the reward is high enough, people tend to choose their number closer and closer to the Nash equilibrium.

From these facts we can conclude that classical game theory does not give an accurate description on how this game is played or should be played to maximize profit. When looking at how humans play the game, it appears that humans tend to assume some degree of cooperation from the other player. Human players realize that cooperating will increase the payout of both players, so a player benefits from choosing a high number, assuming the other player chooses a high number as well. Even game theory experts, who should know about the Nash equilibrium and the theoretical problems in this dilemma, tend to assume a degree of cooperation from the other player.

In short, the 'optimal' strategy provided by classical game theory does not give a satisfying or realistic results, but looking at computer simulations or at humans playing the game can provide better explanations on how the game is played, or should be played, in practice.

Bibliography

[1] Kaushik Basu *The Traveler's Dilemma*, Scientific American, June 2007.
http://www.sciam.com/article.cfm?articleId=7750A576-E7F2-99DF-3824E0B1C2540D47

[2] Kaushik Basu *The Traveler's Dilemma: Paradoxes of Rationality in Game Theory*, The American Economic Review (1994), Vol. 84, No. 2, pp. 391-395

[3] Michael H. F. Wilkinson *Game Theoretical Approaches to Modelling and Simulation (Lecture slides)*, Institute for Mathematics and Computing Science, University of Groningen, November 2004.

[4] J. Maynard Smith *The Theory of Games and the Evolution of Animal Conflicts*, Journal of Theoretical Biology (1974) Vol. 47, pp. 209-21

[5] B. Thomas *On evolutionarily stable sets* J. Math. Biology (1985) Vol. 22 pp. 105-115.

[6] Daniel Ashlock *Evolutionary Computation for Modeling and Optimization*, Springer (2005)

[7] charlie: A genetic algorithms library for Ruby
http://rubyforge.org/projects/charlie/

[8] Rineke Verbrugge and Lisette Mol *Learning to apply Theory of Mind*, January 2008, to appear in *Journal of Logic, Language and Information.*

[9] T. Becker, M. Carter and J. Naeve *Experts playing the Traveler's Dilemma*, Universität Hohenheim (2005)

[10] C. M. Capra, J. K. Goeree, R. Gomez, C. A. Holt *Anomalous Behavior in a Traveler's Dilemma?*, The American Economic Review (June, 1999), Vol. 89, No. 3, pp. 678-690.

Appendix A

Simulation Code

All of these programs require a Ruby interpreter and the charlie library. The latter can be installed by using the 'gem' tool included with most Ruby distributions (`gem install charlie`).

Test tool with graphic user interface. Population is displayed on screen and stats (minimum, maximum, average and standard deviation) are written to `stats_history`. Ruby/TK required (included in the windows 'one-click' installer, for linux use the package `libtcltk-ruby1.8`).

```
%w[rubygems charlie tk].each{|x|require x}

class traveler < FloatListGenotype(1,2..100)
  BONUS = 2
  def choice
    genes[0].between(2,100).round
  end
  def fight_points(other)
    c,oc = choice, other.choice
    c == oc ? [c,c] : c < oc ? [c+BONUS,c-BONUS] : [oc-BONUS,oc+BONUS]
  end
  use NullCrossover
end


BNS,PRB,SZ,GRP,GENS=*0..4
lbl = [["Bonus",2],["Mutation Probability",0.5],["Mutation Size (sigma)",5],["Selection Tournament Size",4],["Generations",1000]]

root  = TkRoot.new() { title "Traveler’s Dilemma Simulation" }
mid,bottom = [0,0].map{ TkFrame.new(root).pack(’side’=>’top’,’fill’=>’x’,’expand’=>true) }

popcv = TkCanvas.new(mid) { width 816; height 100; }.pack
entries = lbl.map{|l,v|
  f = TkFrame.new(bottom).pack(’side’=>’top’,’fill’=>’x’)
  TkLabel.new(f) {text l }.pack(’side’=>’left’)
  TkEntry.new(f) { set v }.pack(’side’=>’right’)
}
$paused = false
$thr = nil
TkButton.new(bottom) {text "[Run]" ;
   command{ $thr.kill rescue nil
    $thr = Thread.new{
     history = []
     traveler::BONUS = entries[BNS].get.to_f
     klass = Class.new(traveler){
       use ListMutator(:probability[entries[PRB].get.to_f],:gaussian[entries[SZ].get.to_f]), CoTournamentSelection(entries[GRP].get.to_i)
       def mutate!; super; genes[0] = genes[0].between(traveler::BONUS,100); end
     }
     pop = Population.new(klass)
     entries[GENS].get.to_i.times{|g| sleep 0.3 while $paused
       gs = pop.evolve_silent(1).map(&:choice).sort
       history << gs.stats
       popcv.delete(’all’)
       TkcLine.new(popcv,8*2,50,800,50)
       gs.each{|x| TkcOval.new(popcv,8*x-5,45,8*x+5,55) }
```

```
      TkcText.new(popcv,400,75,:text=>"Generation #{g} : " + gs.map{|x|x.to_s}.join(', '))
    }
    File.open('stats_history','w'){|f| history.transpose.each{|a| f << a.join(' ') << "\n" } }
   }
  }
}.pack
TkButton.new(bottom) {text "[Paused]" ;  command{ $paused = !$paused } }.pack
Tk.mainloop
```

Program that generated most of our results. When given the reward as a parameter, runs the algorithm for a variety of mutation and tournament size settings and outputs the average strategy in the population after 1000 generations. Ruby/TK not required.

```
require 'rubygems'
require 'charlie'

exit(puts "GIVE ARG!") unless ARGV[0]

BONUS = ARGV[0].to_i
MIN = BONUS
MAX = 100
class traveler < FloatListGenotype(1,MIN..MAX)
  def choice
    genes[0].between(MIN,MAX).round
  end
  def fight_points(other)
    c,oc = choice, other.choice
    c == oc ? [c,c] : c < oc ? [c+BONUS,c-BONUS] : [oc-BONUS,oc+BONUS]
  end
end


for group_size in (2..20)
  $stderr.puts group_size
  for mut_size in (1..100).map{|x| x / 10.0 }
    klass = Class.new(traveler){
      use ListMutator(:probability[0.5],:gaussian[mut_size]), CoTournamentSelection(group_size)
      def mutate!; super; genes[0] = genes[0].between(MIN,MAX); end
    }
    print Population.new(klass).evolve_silent(1000).map{|x| x.genes[0] }.average, ' '
  end
  puts
end
```