

Basu, Kaushik; Becchetti, Leonardo; Stanca, Luca

Working Paper

Experiments with the traveler's dilemma: Welfare, strategic choice and implicit collusion

CAE Working Paper, No. 08-07

Provided in Cooperation with:

Center for Analytical Economics (CAE), Cornell University

Suggested Citation: Basu, Kaushik; Becchetti, Leonardo; Stanca, Luca (2008) : Experiments with the traveler's dilemma: Welfare, strategic choice and implicit collusion, CAE Working Paper, No. 08-07, Cornell University, Center for Analytical Economics (CAE), Ithaca, NY

This Version is available at:

<https://hdl.handle.net/10419/70437>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

CAE Working Paper #08-07

**Experiments with the Traveler's Dilemma:
Welfare, Strategic Choice and Implicit Collusion**

by

Kaushik Basu
Leonardo Becchetti
and
Luca Stanca

October 2008

Experiments with the Traveler's Dilemma: Welfare, Strategic Choice and Implicit Collusion*

Kaushik Basu[†], Leonardo Becchetti[‡], Luca Stanca[§]

October 31, 2008

Abstract

This paper investigates behavior in the Traveler's Dilemma game and isolates deviations from textbook predictions caused by differences in welfare perceptions and strategic miscalculations. It presents the results of an experimental analysis based on a 2x2 design where the own and the other subject's bonus-penalty parameters are changed independently. We find that the change in own bonus-penalty alone entirely explains the effect on claims of a simultaneous change in one's own and the other's bonus-penalty. An increase in the other subject's bonus-penalty has a significant negative effect on claims when the own bonus-penalty is low, whereas it does not have a significant effect when the own bonus-penalty is high. We also find that expected claims are inconsistent with actual claims in the asymmetric treatments. Focusing on reported strategies, we document substantial heterogeneity and show that changes in choices across treatments are largely explained by risk aversion.

*We thank participants at the ESA 2008 and IAREP 2008 annual meetings for their useful comments and suggestions, and Monica Capra for kindly providing the instructions for the Capra et al. (2004) experiment. We also thank Tommaso Reggiani for excellent research assistance. Financial support from Econometica is acknowledged.

[†]Department of Economics, Cornell University; Ithaca, NY, 14853.

[‡]Department of Economics, University of Rome - Tor Vergata.

[§]Economics Department, University of Milan Bicocca. Piazza dell'Ateneo Nuovo 1, 20126 Milan, Italy. E-mail: luca.stanca@unimib.it

1 Introduction

There is now a substantial body of laboratory evidence that human beings do not play games in the way that game theory, founded on the assumption of individual rationality, predicts they will. Unlike in non-strategic, decision-theoretic contexts, a violation of “rational play” in *game-theoretic* contexts can be dissected into two broad categories: (1) a deliberate use on the part of a player of non-selfish and pro-(or, for that matter anti)-social considerations, and (2) a failure to do one’s *strategic* calculations correctly. Laboratory experiments with choice in games, as opposed to ordinary decision-making, allows us to dissect between the above two reasons for deviation from standard theory. This paper reports on a set of experiments designed by us especially to differentiate between these two reasons for deviation.

Such an exercise has important implications for policymaking and the analysis of welfare. The non-uniqueness of the relation between perceptions of welfare and actions chosen was analysed in a celebrated paper by Prasanta Pattanaik (1968) and is also the subject of several papers by Sen (see, in particular, Sen, 1977). One of the earliest works demonstrating the close connection between game-theoretic decision-making and group welfare was Pattanaik (1978). If human deviations from the predictions of standard economics were caused entirely by the fact of a player treating the payoffs differently from the ones represented in the game because of social, altruistic or other such considerations, there would be less of a case for third party intervention than if the deviations were caused by a systematic failure to do one’s strategic calculations right. In the former case we would simply have to admit that an individual’s own perception of his or her welfare is different from that of the analyst’s, and most analysts in such contexts would be willing to have the analyst’s perception of welfare be over-ruled by the actual individual’s perception of his or her welfare. But that would not be so compelling in contexts where individuals could be shown to be demonstrably prone to ignoring relevant strategic considerations.

The game that we use to study the above problem is the Traveler’s Dilemma (Basu, 1994, 2007). This game (henceforth, TD) is an example of a strategic setting where the assumption of rationality commonly made in mainstream economics produces counter-intuitive predictions, both because of the complexity of the strategic decision analysis and because pro-social preferences can pull one away from selfish behavior. Before discussing this, let us briefly recall the game. In the original version of the TD, two players

must individually choose an integer between 2 and 100, without communicating with each other. If both players choose the same number, they get a payoff equal to that number. If they choose different numbers, the player who chooses the lower number gets that number plus a bonus (+2, in the original story), while the player who chooses the higher number gets the lower number minus a penalty (-2 , in the original story). It is easy to verify that the only Nash equilibrium in this game is where both players make the minimum claim of 2. This is true even when mixed strategies are allowed. Further, (2, 2) is also the only trembling hand perfect equilibrium, the only strict equilibrium, and the only rationalizable outcome. Yet most people on introspection feel that they would play it differently. Moreover, when the game has actually been tested experimentally, people consistently reject the Nash equilibrium choice (Capra et al., 1999; Becker, Carter and Naeve, 2005; Rubinstein, 2006, 2007; Chakravarty, Dechenaux and Roy, 2008). Indeed, by rejecting the rational choice, agents end up obtaining larger rewards.

Capra et al. (1999) find that, in contrast to theoretical predictions, the size of the bonus-penalty matters: if it is small, claims converge to the maximum claim; if it is large, claims converge to Nash Equilibrium. Becker et al. (2005) show that, even when the TD is played by experts, so that ignorance of (introspective) backward induction can be ruled out, average claims are much higher than the Nash equilibrium. Expected claims are also higher than in the Nash equilibrium, so that it is rational not to play the Nash equilibrium outcome. This seems to suggest the problem lies with the iterated use of rationality that is entailed by rationality being common knowledge. As this may suggest, the TD raises both an experimental and a theoretical question. The latter asks the following: Even if both players were fully rational and this was common knowledge, does this have to imply that the outcome will be the Nash equilibrium? This theoretical question has also given rise to a lot of discussion and controversy and remains largely unresolved (see Basu, 2000; Colombo, 2003; Zambrano, 2004; Brandenburger, 2007). Much of this theoretical literature is predicated on the assumption that players choose strategically not to play the Nash strategy and that such deviations, despite appearing otherwise, are rational.

The question we address in the present paper is however the practical one: How do ordinary individuals actually choose? We present an experimental analysis based on a novel 2x2 design, where the own and the other subject's bonus-penalty parameters are changed independently, either symmetrically, as in the standard TD, or asymmetrically. This allows us to shed light on the

result that the size of the theoretically irrelevant bonus-penalty matters in practice. More generally, it allows us to explore the determinants of strategic behavior in a one-shot TD and separate out strategic and altruistic reasons for deviating from Nash behavior.

We find that a joint change in own and other subject’s bonus-penalty has a large and significant effect on claims, thus extending to a one-shot setting the result in Capra et al. (1999) for a repeated game. A change in the other subject’s bonus-penalty has a significant but relatively small effect on claims when the own bonus-penalty is low, while it has no net effect when the own bonus-penalty is high. On the contrary, a change in own bonus-penalty has a large and significant effect on claims, for both low and high level of opponent’s bonus-penalty. Comparing directly the two asymmetric treatments, a change in own bonus-penalty has a significantly larger effect on claims than a change in other’s bonus-penalty. This suggests that players do not take full account of strategic considerations. We also examine subjects’ beliefs and their reported strategies, as provided in a post-experimental questionnaire. We find that, in the asymmetric treatments, expected claims are not consistent with actual claims. Finally, focusing on reported strategies, our findings document that changes across treatments are also driven by risk aversion.

The paper is structured as follows. Section 2 describes the experimental design, the hypotheses to be tested and the procedures. Section 3 presents the results at the aggregate level. Section 4 provides an interpretation of the results, examining the data at the individual level, subjects’ beliefs and reported strategies. Section 5 concludes with a discussion of the main findings and the implications of the analysis.

2 Experimental design and procedures

The baseline game is a standard TD with the same set of parameters as in Capra et al. (1999), in order to ease comparability. Two subjects must individually choose an integer between 80 and 200, without communicating with each other. If both subjects make the same claim, they will each be paid that amount. If they make different claims, the player who makes the lower claim obtains that amount plus a bonus, while the player who makes the higher claim obtains the lower amount minus a penalty.

The experiment is based on a design in which the treatment variables

are the *own* and the *other subject's* bonus-penalty parameters (henceforth indicated by R and \tilde{R} , respectively). We vary the two treatment variables independently, setting their values at 10 and 80, thus obtaining four treatments in a 2x2 design, as described in table 1.¹ This design allows us to distinguish between the direct (net) effects of the treatment variables (R and \tilde{R}) and their interactive effects.

Table 1 about here

2.1 Hypotheses

Under the assumptions that all agents are rational, and that rationality is common knowledge, the theoretical prediction, according to the familiar backward induction argument (see Priest, 2000), would be that both players choose the minimum feasible number, namely, 80, in all treatments. The exact size of both R and \tilde{R} is irrelevant, as long as they are greater than 1. If we drop these assumptions of full rationality or, more minimally, the common knowledge of full rationality, independent changes in R and \tilde{R} allow to distinguish between conditional and unconditional behavior. In the former case, agents try to formulate an expectation about the other player's claim and maximize their expected payoff accordingly. Their claims can therefore be affected by changes in *either* R or \tilde{R} . In the case where the agents disregard strategic thinking, they maximize their own payoff taking as given the other player's strategy, unmindful of the fact that a change in \tilde{R} can alter the other player's behavior. Their claims can be affected only by changes in R , whereas changes in \tilde{R} should be irrelevant.

Defining μ_i as the mean claim in treatment i , we test the following hypotheses:

H1. Effect of a change in both *own* and *other subject's* bonus-penalty:

$$H_0 : \mu_2 = \mu_1 \text{ vs } H_1 : \mu_2 < \mu_1 \quad (1)$$

The null hypothesis is the irrelevance of a joint increase in R and \tilde{R} , versus the alternative hypothesis of a negative effect on claims. This hypothesis, rejected by Capra et al. (2004) in a similar setting with

¹Note that treatments 3 and 4 are identical between subjects, whereas they are different within subjects.

repeated interaction, is tested here in a one-shot setting and is used as a benchmark.

- H2. Effect of a change in *other subject's* payoff only. The null hypothesis is the irrelevance of an increase in \tilde{R} , versus the alternative hypothesis of a negative effect on claims. This hypothesis can be tested under two scenarios for the own bonus-penalty:

H2a. Keeping constant *low* R :

$$H_0 : \mu_3 = \mu_1 \text{ vs } H_1 : \mu_3 < \mu_1 \quad (2)$$

This test assesses the gross effect of an increase in \tilde{R} . Rejection of the null hypothesis provides an indication of conditional behavior.

H2b. Keeping constant *high* R :

$$H_0 : \mu_2 = \mu_4 \text{ vs } H_1 : \mu_2 < \mu_4 \quad (3)$$

This test evaluates the net effect of an increase in \tilde{R} , allowing for the interaction with R .

- H3. Effect of change in *own* bonus-penalty only. The null hypothesis is the irrelevance of an increase in R , versus the alternative hypothesis of a negative effect on claims. This hypothesis can be tested under two scenarios for the other subject's bonus-penalty:

H3a. Keeping constant *low* \tilde{R}

$$H_0 : \mu_4 = \mu_1 \text{ vs } H_1 : \mu_4 < \mu_1 \quad (4)$$

This test assesses the gross effect of an increase in R . Rejection of the null hypothesis is consistent with either conditional or unconditional behavior.

H3b. Keeping constant *high* \tilde{R}

$$H_0 : \mu_2 = \mu_3 \text{ vs } H_1 : \mu_2 < \mu_3 \quad (5)$$

This test examines the net effect of an increase in R , allowing for the interaction with \tilde{R} .

H4. Effect of change in own versus other subject's bonus-penalty:

$$H_0 : \mu_4 = \mu_3 \text{ vs } H_1 : \mu_4 \neq \mu_3 \quad (6)$$

The null hypothesis is that changes in R and \tilde{R} have the same effect on claims versus the alternative of different effects. This test provides a direct comparison of the effects of the own and the other subject's bonus-penalty.

2.2 Procedures

We ran four sessions, with 24 subjects participating in each session, for a total of 96 subjects. We used a within-subjects design, so that in every session each subject played the four treatments in four sequential phases. Subjects knew that only one phase would be drawn randomly to determine payoffs. The effect of repetition was controlled for by the randomization of treatments within sessions. Subjects were informed that they would never interact more than once with the same subject, in order to avoid strategic incentives. In addition, subjects only received feed-back about the outcomes of each phase at the end of the four phases, in order to minimize the effects of learning and avoid cross-subject dependence.

Beliefs about opponents' claim were elicited as a surprise question after the implementation of the four treatments. Subjects could win 5 euros by correctly guessing the other subject's claim for the selected treatment. A surprise was necessary so as not to impose subjects to think about what they expected their opponent to play. We used a point-expectation, rather than an interval or a distribution, in order to avoid strategic responses.

In each of the four sessions, subjects were randomly assigned to a computer terminal at their arrival. In order to ensure public knowledge, instructions were distributed and read out aloud (see Appendix 1). Sample questions were distributed to ensure understanding of the experimental instructions. Answers were privately checked and, if necessary, explained to the subjects. The experiment did not start until all subjects had answered all questions correctly.

The experiment was conducted in the Experimental Economics Laboratory of the University of Milan Bicocca in April 2008. Participants were undergraduate students of Economics recruited by e-mail using a list of voluntary potential candidates. Sessions lasted approximately 45 minutes. No

show-up fee was paid and the exchange rate was 10 points = 1 euro. Theoretical payments ranged between 0 and 33 euros, actual payments were between 0 and 27 euro, with an average of about 12 euros. The experiment was run using the experimental software z-Tree (Fischbacher, 2007).

3 Results

Figure 1 displays mean and median claims and expected claims for each treatment across all phases. Average claims are 174.4 and 115.6 in the 10-10 and 80-80 treatments, respectively. Median claims are 199 and 80, respectively. Average claims are 112.3 and 149.1, respectively, in the 80-10 and 10-80 treatments, while median claims are 80 and 159, respectively. This preliminary description indicates that a simultaneous change in both own and other's bonus has a large effect on claims even in a one-shot setting. However, in the asymmetric treatments, a change in other's bonus-penalty alone has a relatively small effect on claims, while a change in own bonus-penalty has a much larger effect. Indeed, the difference between the symmetric treatments (10-10 and 80-80) can be largely explained by the change in the own bonus-penalty.

Figure 1 about here

Table 2 reports the corresponding figures mean and median claims by individual phases and overall. The table indicates, for all treatments, a tendency for claims to fall over successive phases, suggesting that learning may be playing a role even in the absence of feed-backs.

Table 2 about here

Table 3 reports results of sign-rank tests of the null hypothesis of equal claims between treatments, based on within-subject matched pairs of observations. The difference between the 10-10 and 80-80 treatments (column 1) is strongly statistically significant, both overall and within individual phases. This result confirms and extends to a one-shot game without repetition the finding of Capra et al. (1999) under repeated interactions and also the recent experimental findings of Chakravarty, Dechenaux and Roy (2008). This experimental result is also consistent with the theory of "iterated regret minimization" of Halpern and Pass (2008).

Result 1: A joint increase in own and other’s bonus-penalty has a large and significant negative effect on claims.

Table 3 about here

Next, consider the effect of a change in the other subject’s bonus, keeping fixed the own bonus (hypotheses 2a and 2b). The difference between the 10-10 and 10-80 treatments (column 2) is positive and statistically significant. On the other hand, comparing the 80-10 and 80-80 treatments (column 5), the difference is not statistically significant. These results indicate that the other subject’s bonus-penalty has a significant gross effect on claims. However, when we remove the interaction with the own bonus-penalty, the net effect is not significant. Note that this finding also implies that the change in the own bonus-penalty by itself entirely explains the effect of a joint change in own and other’s bonus-penalty.

Result 2: An increase in the other subject’s bonus-penalty has a significant negative effect on claims when the own bonus-penalty is low, whereas it does not have a significant effect when the own bonus-penalty is high.

The third set of hypotheses refers to the effect of a change in own bonus, for a given level of the other subject’s bonus. The difference between the 10-10 and 80-10 treatments (column 3) is strongly statistically significant, not only overall but also by individual phases. Comparing the 10-80 and 80-80 treatments (column 4), the difference is also statistically significant. This indicates that, contrary to the other subject’s bonus-penalty, the own bonus-penalty does have a significant *net* impact on claims.

Result 3: An increase in the own bonus-penalty has a large and significant negative effect on claims when the other subject’s bonus-penalty is low and a significant negative effect also when the other subject’s bonus-penalty is high.

Finally, when comparing the two asymmetric treatments (column 6), the difference between the 10-80 and 80-10 treatments is negative and statistically significant.

Result 4: A change in one’s own bonus-penalty has a significantly larger effect on claims than a change in the other’s bonus-penalty.

4 Analysis

This section explores what lies behind the results at the aggregate level presented above. We start by examining agents' beliefs about other subjects' claims. We then focus on choices and revealed strategies at the individual level. Finally, we examine the strategies reported by subjects in a post-experimental questionnaire.

4.1 Beliefs

Table 4 reports mean and median expected claims for each treatment, by individual phase and overall. Mean and median beliefs by treatment across all phases are also displayed in the bottom panels of Figure 1. Average expected claims in the 10-10 and 80-80 treatments are 172 and 110.4, respectively, while median expected claims are 200 and 80, respectively: the change in both own and other's bonus has a large effect on expected claims, consistently with the effect on claims.

Table 4 about here

The results for the asymmetric treatments are quite surprising. In the 80-10 and 10-80 treatments, average expected claims are 121.4 and 135.9, respectively (median expected claims are 80 and 130, respectively). If players had rational expectations, beliefs could be expected to be lower when the other player's bonus-penalty were high. What we observe is indeed the opposite: expected claims are higher (lower) when the other player's bonus-penalty is high (low) and the own bonus-penalty is low (high). In asymmetric treatments, players do not seem to be able to disentangle their own from the other player's expected behaviour. Hence, this suggests that the deviation from the Nash outcome is prompted not by considerations of altruism and pro-social behavior but by an inability to do strategic analysis and take into consideration the other player's behavioral response.

Result 5: In the asymmetric treatments, beliefs are not consistent with choices.

This finding does not seem to be due to a general inaccuracy in formulating expected claims since, as indicated in the next section, the distribution of the prediction error reveals that expected claims are on average, and in larger number, correct.

4.2 Individual Choices

Figure 2 displays the distribution of individual claims and expected claims, by treatments and overall. In order to shed light on subjects' strategies, we also match actions and beliefs at individual level. Table 5 reports cross-tabulations for claims (rows) and expected claims (columns), by treatment and overall. Across the four treatments, 38 per cent of the subjects make the minimum claim of 80, 25 per cent the maximum claim of 200, while 37 per cent fall within the 81-199 range. As for expected claims, 43 per cent of players believe that their counterpart will choose the NE claim, while in 34 per cent of the cases the expected claim is 200. The NE combination for claim and expected claim (80,80) occurs in 32 percent of cases, while the (200,200) combination occurs in 21 per cent of the cases.

Figure 2 about here

Table 5 about here

When comparing individual treatments, subjects playing 80 and 200 are 8 and 44 per cent, respectively, in the 10-10 treatment, while they are 57 and 13 per cent, respectively, in the 80-80 treatment. The change is indeed even more pronounced in the 80-10 treatment, where subjects claiming 80 and 200 are 61 and 9 per cent, respectively. It is interesting to observe that the distributions for the 80-10 and 80-80 treatments are virtually identical. This confirms the finding that the observed change in claims between the 10-10 and 80-80 treatments is entirely attributable to the change in the own bonus.

Table 6 reports the cross tabulation of observed frequencies obtained by comparing the 10-10 benchmark treatment with each of the other three treatments, hence providing information about changes in claims across treatments within subjects. The results indicate a tendency towards polarization, as 25 per cent of the subjects playing 200 in the 10-10 treatment, switched to 80 in the 80-10 treatment. Only 11 per cent of the subjects played 200 in both the 10-10 and 80-80 treatments. An additional 28 per cent of the subjects played between 81 and 199 in the 10-10 treatments and switched to 80 in the 80-10 treatment.

Table 6 about here

Figure 3 displays the distribution of prediction errors, defined as the difference between the own belief and the other subject's choice. Across all

treatments, beliefs are correct for about 30 per cent of the subjects. Within treatments, other subjects' claims tend to be overpredicted in the 10-80 treatment and, conversely, underpredicted in the 80-10 treatment. This finding further illustrates the result that in the asymmetric treatments players do not disentangle their belief on the other subject's choice from their own behaviour (see also the surprisingly similar shape of claims and expected claims in the asymmetric treatments in Figure 1). One possible interpretation of this result is that, in a number of cases, agents make choices without formulating an expectation about other subjects' claim. When beliefs are elicited, *ex post*, agents use their decisions to formulate their beliefs about other subjects' claims. For an overall evaluation of the reliability of beliefs we must put together the two above mentioned conflicting facts. On the one hand, beliefs correct in 30 per cent of the cases and the symmetric distribution of prediction error seem to indicate the accuracy of formulated beliefs. On the other hand, the overprediction and underprediction described above document a cognitive bias in the asymmetric treatments.

Figure 3 about here

4.3 Revealed strategies

Table 7 reports the cross tabulation of observed frequencies for subject types. We identified subject types on the basis of claims and beliefs as follows. Subject playing and expecting 80 are defined NASH. Subjects playing slightly less than the expected claim (between 1 and 5 units) are classified as strategic (STRA).² Subjects claiming the same amount they expect (except for 80,80) are classified as team strategic (TEAM). Subjects who claim less than 5 units than expected are defined weakly rational (WEAK), and those who claim more than the expected bid are defined irrational (IRRA). In general, we observe that the increase in players' bonus-penalty reduces the violation of individual rationality or team rational choices and increases Nash rational outcomes. For example, the (80,80) NE pair occurs only in 4 per cent of cases in the 10-10 treatment, while in 53 per cent of cases in the 80-80 treatment. The already evidenced dominance of the player's penalty change over the counterpart's penalty change is supported by the fact that the claim-belief

²By reasonably assuming that the distributions of the expected claims are non degenerate, players may increase the probability of winning the reward by choosing $C < C^*$, where $C^* = C^e - 1$. The -5 threshold is obviously arbitrary.

pair is (80,80) in 22 per cent of cases in the 10-80 treatment, while in 49 per cent of cases in the 80-10 treatment (close to what happens in the 80-80 one). On the other hand, the team-strategic consistent pair (200, 200) occurs in 38 per cent of cases in the 10-10 treatment and in 10 per cent of cases in the 80-80 treatment. Finally, irrational strategies tend to fall when we move from the 10-10 treatment (13 per cent) to treatments with high bonus-penalty (only 1 per cent in the 80-80 treatment). Some subjects appear to play randomly when consequences are not severe, while they tend to concentrate and use rationality when monetary consequences are more serious.

Table 7 about here

4.4 Reported Strategies

In this section we examine subjects' ex-post descriptions of their strategies. Although it is difficult to classify unequivocally self-reported descriptions, three types of strategies are clearly identifiable:

1. Conditional: formulate an expectation about the other player's claim, and play accordingly.
2. Risk-averse: try to minimize possible loss, focusing on the size of the penalty.
3. Collusive (risk-loving): play high hoping the other also plays high.

Figure 4 displays the distribution of the reported strategies. Players emphasizing conditional behavior are about 17 per cent. Strategies based on collusion and risk aversion are about 17 and 29 per cent, respectively. A further 33 per cent declares other strategies which are not clearly classifiable, while about 4 per cent does not declare any strategy. Note that conditional players in many cases just say that they take into account their expectation on the counterpart's action, while in other cases they add that they will try to undercut them.

Figure 4 about here

Figure 5 displays average claims by declared strategy type for each of the four treatments. Figure 6 reports average changes between treatments by declared strategy type. Subjects whose declared strategies are based on risk aversion display the largest negative change between treatments when the

own bonus-penalty is increased. Subjects who reported conditional strategies display the largest effect on claims, as expected, when the other subject's bonus-penalty is increased. Finally, it is interesting to observe that the claims of subjects who reported collusive strategies are relatively unaffected when the own and other subject's bonus-penalty are jointly increased.

Figures 5 and 6 about here

In order to assess the statistical significance of these treatment effects by reported strategy type, table 8 reports OLS estimates obtained when the within-subject change between pairs of treatments is regressed against a constant (overall) or against a set of dummy variables for subject types identified on the basis of the reported strategies described above (by strategy), in order to identify the contribution of individual types to the overall change. The dependent variable in each column is the difference between the claims in the two treatments indicated in the column headings.

Table 8 about here

A relevant finding is the large and significant negative coefficient of the dummy for declared risk aversion in all specifications, with the only exception of the difference between the 80-80 and 10-80 treatments. Risk aversion is indeed not only the most commonly reported determinant of choices (28 per cent), but also the strategy associated with the largest effects of treatment variables on claims.

Result 6: Treatment effects on claims are largely explained by risk aversion.

This result is consistent with a standard mean-variance payoff utility function. A higher bonus-penalty structure expands the payoff range and increases its variance. As a consequence, the reduction of claims will be higher for individuals with higher risk aversion. Note as well that risk aversion, combined with non-Nash rationality, may explain the sensitivity of subjects' claims to changes in the bonus-penalty structure.

The results in table 8 also indicate that the positive change in claims produced by an increase in the other's subject bonus-penalty, keeping the own bonus-penalty high, is explained by the behavior of collusive players. More generally, it can be observed that collusive players tend to claim less when the structure of bonus-penalties is asymmetric.

Finally, in order to provide a robustness check, table 9 reports Tobit estimation results for the same set of specifications, to take into account the truncated nature of the dependent variables. All the results described above are qualitatively unchanged.

Table 9 about here

5 Concluding Remarks

The Traveler’s Dilemma and other related games, such as the Centipede and the finitely-repeated Prisoner’s Dilemma, suggest that the Nash equilibrium predictions do not work in contexts where the rationality assumptions are too demanding and rely on higher order knowledge of the rationality of players. Once it is experimentally verified that players frequently deviate from Nash equilibrium, the natural question that arises is: How do they actually choose? Are the deviations systematic and can we parse the deviations to gain insights into individual motivations and cognition?

In order to shed light on the blackbox of the decision process of ordinary human beings, in this paper we pursued three original directions. (i) We decomposed the bonus-penalty change of Capra et al. (1999) into its two asymmetric change components. (ii) We used a design under which the same subject plays different treatments without learning about previous outcomes. An added value of this approach is its closeness to the first best of the comparison of a treatment with the counterfactual: players are subject to different treatments almost at the same time (in immediately subsequent phases) without any feedback on previous plays. (iii) We collected players’ expected claims and ex post declared strategies and thereby made it possible to compare these with their actual plays.

The most relevant findings of our research can be clustered under four categories. First, the dominance of the change in one’s own bonus-penalty over the change in the other player’s bonus-penalty is such that the former explains almost all the experimental results of Capra et al. (1999). Hence, our experiments, taking cue from the work of Capra et al., help isolate and parse more proximate causes of what prompts deviation from Nash behavior. This is confirmed in many ways (descriptive evidence in mean and median claims and expected claims, transition across different treatments, direct nonparametric tests, etc.).

Second, heterogeneity of players' preferences is supported by observed claim-belief pairs across different treatments and self-revealed strategies at the end of the game. We interpret this variability in terms of three different motivational types—Nash or individually rational, team strategic and irrational. Third, even though 30 per cent of expected claims are correct and the distribution of the prediction error around the zero mean is symmetric (suggesting that expected claims were in general carefully formulated), we document a cognitive bias in asymmetric treatments where claims tend to be underpredicted in the 80-10 treatment and, conversely, overpredicted in the 10-80 treatment. More simply, it seems that players are not able to distinguish fully between their own and their opponent's behaviour, which suggests an inherent inadequacy in strategic thinking. Fourth, reported strategies help to explain changes in claims when the penalty-reward structure varies with respect to the 10-10 benchmark. More specifically, players we classify as risk averse tend to claim significantly less and those classified as collusive do the same but only in asymmetric treatments.

References

- Barker, A. (2008) Evolutionary Stability in the Traveler's Dilemma, College Mathematics Journal, forthcoming.
- Basu, K. (1994) The Traveler's Dilemma: Paradoxes of Rationality in Game Theory, American Economic Review, Vol. 84, No. 2, 391-395.
- Basu, K. (2000), Prelude to Political Economy: A Study of the Social and Political Foundations of Economics, Oxford: Oxford University Press.
- Basu, K. (2007), The Traveler's Dilemma, Scientific American, June, Vol. 296: 90-95.
- Brandenburger, A. (2007), The Games We Assay, Scientific American, Letters, October, Vol. 296: 14.
- Becker, T., Carter, M., and Naeve, J. (2005) Experts Playing the Travelers Dilemma, Working Paper 252, Institute for Economics, Hohenheim University.
- Cabrera, S., Capra, M. C., Gomez, R. (2006) Behavior in One-shot Traveler's Dilemma Games: Model and Experiment with Advice, Spanish Economic Review, Vol. 9, No. 2: 129-152.
- Capra, M., Cabrera, S., Gomez, R. (2003) The Effect of Common Advice on One-shot Traveler's Dilemma Games: Explaining Behavior through an Introspective Model with Errors, CentrA working paper n. 2003/17.
- Capra, M., Cabrera, S., Gomez, R. (2004) Introspection in one-shot traveler's dilemma games, Department of Economics, Emory University, working paper 5/2004.
- Capra, M., Goeree, J.K., Gomez, R. and Holt, C.A. (1999) Anomalous Behavior in a Traveler's Dilemma? American Economic Review, Vol. 89, No. 3: 678-690.
- Chakravarty, S., Dechenaux, E. and Roy, J. (2008) Imprecise and Precise Communication in the Traveler's Dilemma, mimeo: Indian Institute of Technology, New Delhi.

- Colombo, F. (2003) The Game Take-or-Play: A Paradox of Rationality in Simultaneous Move Games, *Bulletin of Economic Research*, Vol. 55, No. 2: 195-202.
- Fischbacher, U. (2007) z-Tree: Zurich Toolbox for Ready-made Economic Experiments. *Experimental Economics*, Vol. 10: 171-178.
- Goeree, J. K. and Holt, C. (2001) Ten Little Treasures of Game Theory, and Ten Intuitive Contradictions, *American Economic Review*, Vol. 91: 1402-1422.
- Halpern, J.Y. and Pass, R. (2008) Iterated Regret Minimization: A More Realistic Solution Concept, mimeo: Cornell University.
- Pattanaik, P. K. (1968) Risk, Impersonability and the Social Welfare Functions, *Journal of Political Economy*, Vol. 76.
- Pattanaik, P. K. (1978), *Strategy and Group Choice*, Amsterdam: North-Holland.
- Priest, G. (2000) The Logic of Backward Inductions, *Economics and Philosophy*, Vol. 16: 267-285.
- Rubinstein, A. (2006) Dilemmas of An Economic Theorist, *Econometrica*, Vol. 74, No. 4, 865-884.
- Rubinstein, A. (2007) Instinctive and Cognitive Reasoning: A Study of Response Times, *Economic Journal*, 117(523), 1243-1259.
- Sen, A. (1977) Rational Fools: A Critique of the Behavioral Foundations of Economic Theory, *Philosophy and Public Affairs*, Vol. 6.
- Zambrano, E. (2004) Counterfactual Reasoning and the Common Knowledge of Rationality in Normal-Form Games, *Topics in Theoretical Economics*, Vol. 4(1), Article 8.

Appendix 1: Instructions

This appendix reports the instructions distributed on paper to the subjects.

Instructions

- Welcome and thanks for participating in this experiment.
- During the experiment you are not allowed to talk or communicate in any way with other participants.
- If at any time you have any questions raise your hand and one of the assistants will come to you to answer it.
- By following the instructions carefully you can earn an amount of money that will depend on your choices and the choices of other participants.
- At the end of the experiment the number of points that you have earned will be converted in euros at the exchange rate $10 \text{ tokens} = 1 \text{ euro}$. The resulting amount will be paid to you in cash.

General rules

- There are 24 subjects participating in this experiment.
- The experiment will take place in 4 independent phases. At the beginning of each phase instructions for that phase will be distributed.
- In each phase 12 couples of two participants will be formed randomly and anonymously, so that in each phase you will interact with a different subject within a couple.
- In each phase the choices that you and the other subject will make will determine the amount earned.
- The choices that you and the other subject will make, and the corresponding results, will not be communicated to you at the end of each phase, but only at the end of the whole experiment.

- At the end of the experiment, **only one** of the four phases will be randomly drawn, and the earnings of each participants will be determined on the basis of the **selected phase**.

PHASE 1,2,3,4

- In this phase you have to choose an integer between 80 and 200.
- At the same time, the subject with whom you have been paired has to choose an integer between 80 and 200.
- If the numbers chosen are the same, you will both earn a number of points equal to the number selected.
- If the numbers chosen are different, you will both earn a number of points equal to the lower of the chosen numbers, plus a **bonus or penalty** determined as follows:
 - [Treatment 1]
 - If the number you have chosen is **smaller** than the number chosen by the other subject, you will have a **bonus of 10 points** and the other subject will have a **penalty of 10 points**.
 - If the number you have chosen is **larger** than the number chosen by the other subject, you will have a **penalty of 10 points** and the other subject will have a **bonus of 10 points**.
 - [Treatment 2]
 - If the number you have chosen is **smaller** than the number chosen by the other subject, you will have a **bonus of 80 points** and the other subject will have a **penalty of 80 points**.
 - If the number you have chosen is **larger** than the number chosen by the other subject, you will have a **penalty of 80 points** and the other subject will have a **bonus of 80 points**.
 - [Treatment 3]
 - If the number you have chosen is **smaller** than the number chosen by the other subject, you will have a **bonus of 10 points** and the other subject will have a **penalty of 80 points**.

- If the number you have chosen is **larger** than the number chosen by the other subject, you will have a **penalty of 10 points** and the other subject will have a **bonus of 80 points**.

[Treatment 4]

- If the number you have chosen is **smaller** than the number chosen by the other subject, you will have a **bonus of 80 points** and the other subject will have a **penalty of 10 points**.
- If the number you have chosen is **larger** than the number chosen by the other subject, you will have a **penalty of 80 points** and the other subject will have a **bonus of 10 points**.

Figure 1: Claims and expected claims: means and medians, by treatment

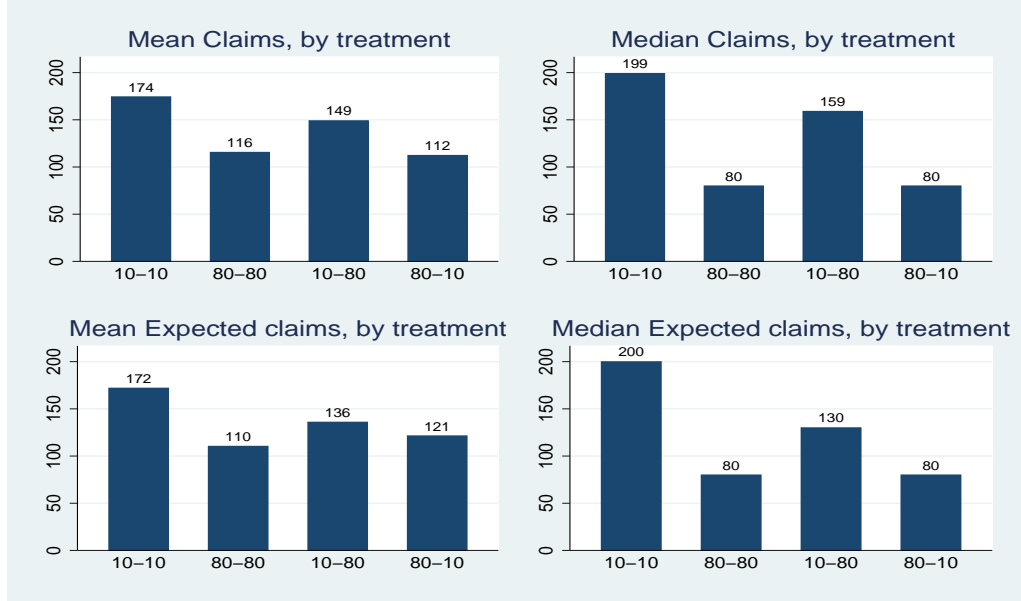


Figure 2: Distribution of claims and expected claims, by treatment

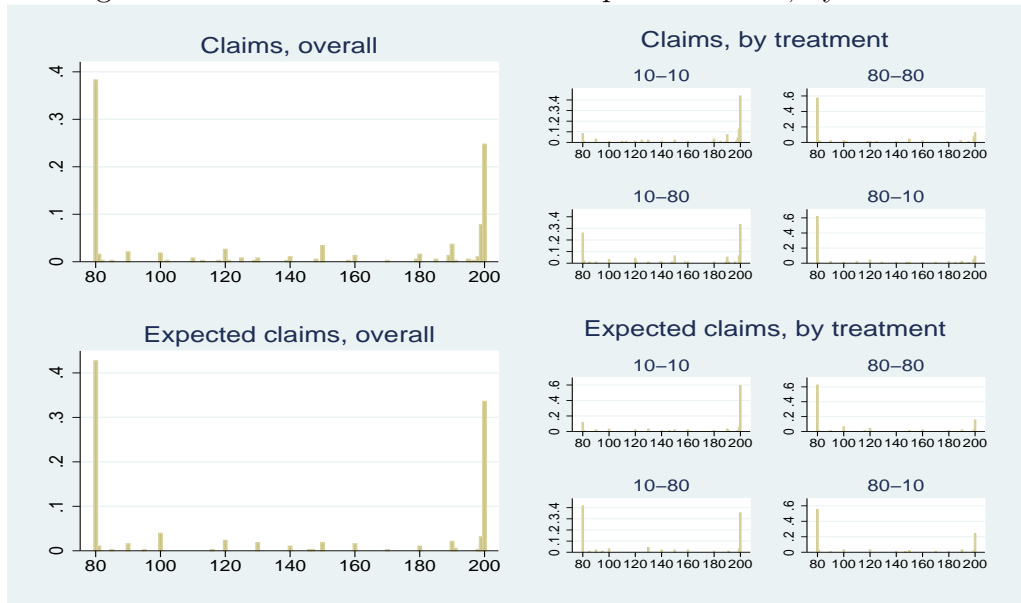


Figure 3: Prediction error, by treatment

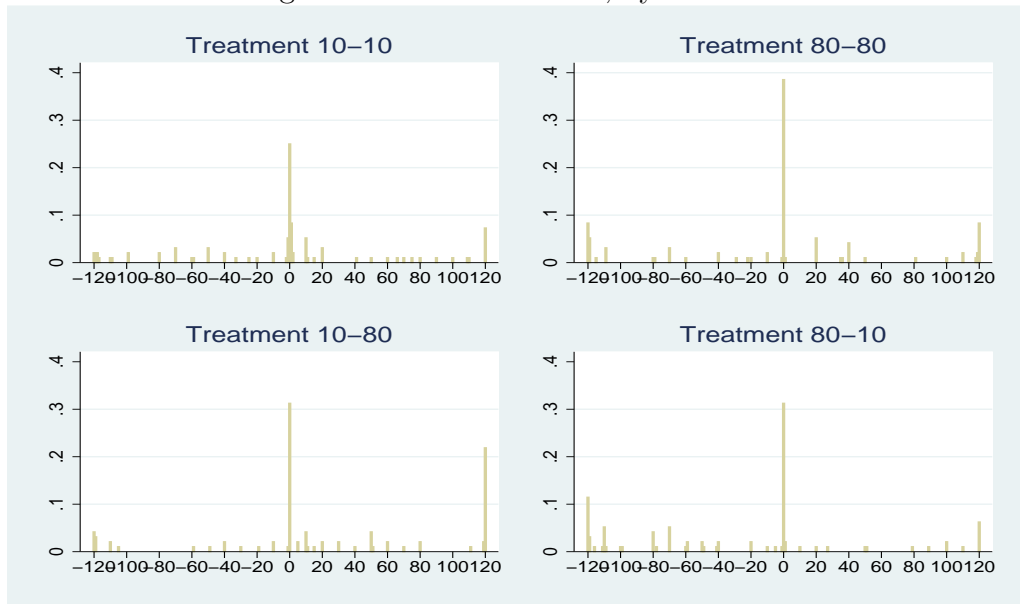


Figure 4: Distribution of reported strategies

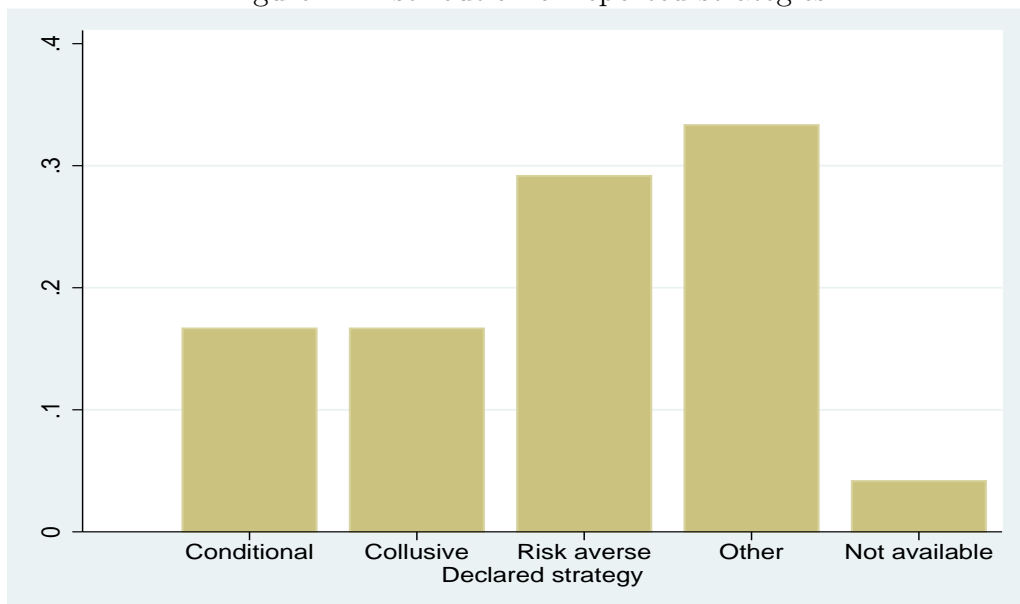


Figure 5: Mean claims, by treatment and strategy type

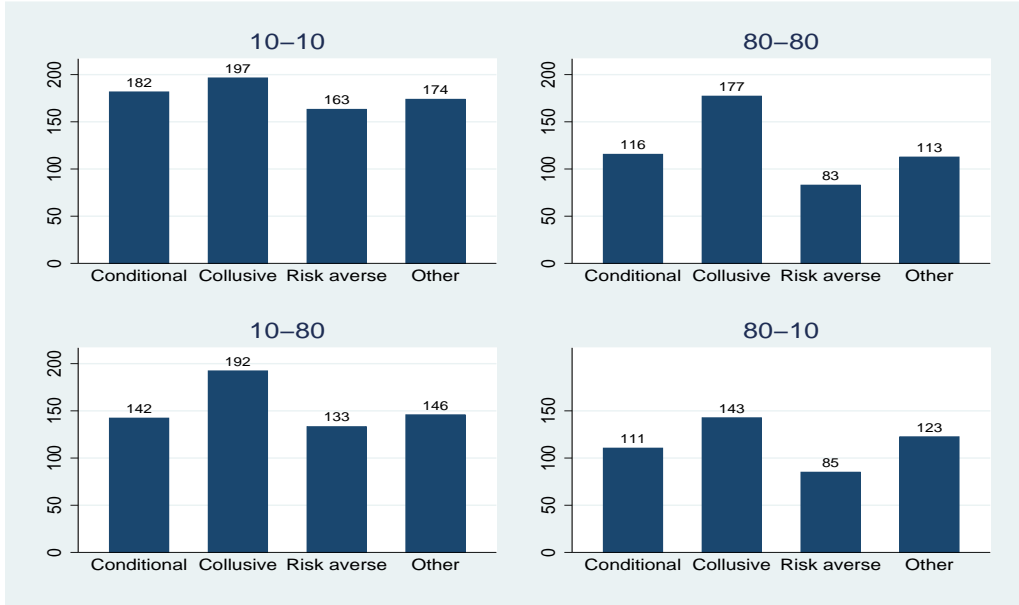


Figure 6: Differences between treatments, by strategy type

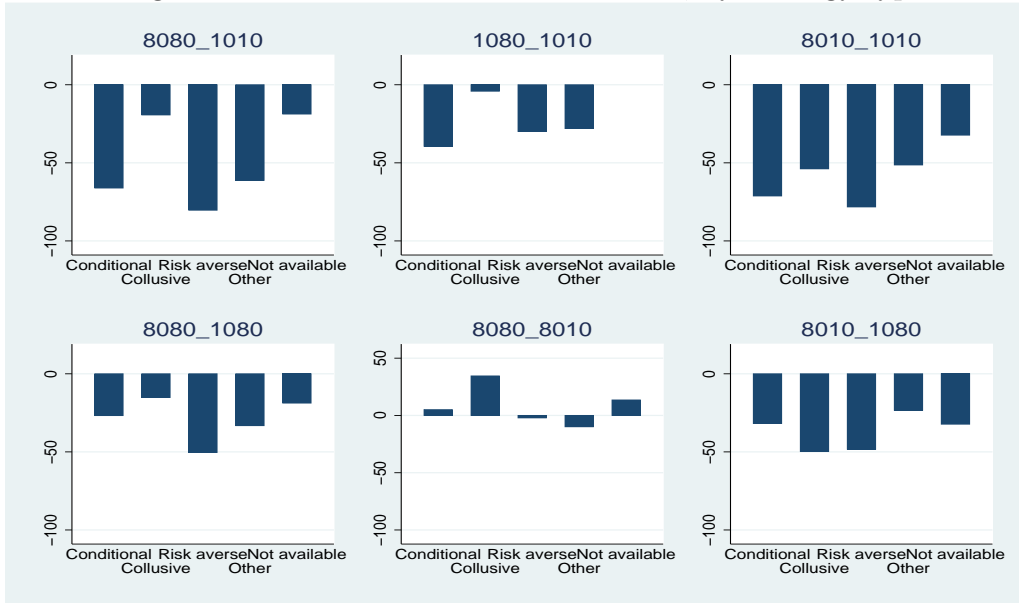


Table 1: Experimental Design

Own bonus-penalty (R)	Other subject's bonus-penalty (\tilde{R})	
	± 10	± 80
± 10	Treatment 1	Treatment 3
± 80	Treatment 4	Treatment 2

Note: see section 2 for details on the experimental design.

Table 2: Mean and median claims, by treatment and phase

	Phase 1	Phase 2	Phase 3	Phase 4	All Phases
<i>Means</i>					
Treatment 1 (10-10)	174.0	176.5	173.4	173.8	174.4
Treatment 2 (80-80)	140.1	111.9	103.8	106.5	115.6
Treatment 3 (10-80)	158.8	139.1	144.2	154.3	149.1
Treatment 4 (80-10)	116.8	116.3	121.5	94.6	112.3
All treatments	147.4	135.9	135.7	132.3	137.8
<i>Medians</i>					
Treatment 1 (10-10)	199.0	199.5	198.5	199.0	199.0
Treatment 2 (80-80)	145.0	80.0	80.0	80.0	80.0
Treatment 3 (10-80)	194.5	150.0	143.5	190.5	159.0
Treatment 4 (80-10)	80.0	80.0	80.0	80.0	80.0
All treatments	164.5	130.0	120.0	90.0	130.0

Note: Treatment 1: $R = \pm 10$ $\tilde{R} = \pm 10$. Treatment 2: $R = \pm 80$ $\tilde{R} = \pm 80$. Treatment 3: $R = \pm 10$ $\tilde{R} = \pm 80$. Treatment 4: $R = \pm 80$ $\tilde{R} = \pm 10$.

Table 3: Tests of equality of claims between treatments, by session and overall

	1010 8080	1010 1080	1010 8010	8080 1080	8080 8010	1080 8010
Session 1	4.03	3.00	3.99	-2.75	-1.16	2.55
(p-value)	0.00	0.00	0.00	0.01	0.24	0.01
Session 2	4.11	0.95	4.10	-3.47	-2.00	3.04
(p-value)	0.00	0.34	0.00	0.00	0.05	0.00
Session 3	3.97	2.24	3.97	-3.24	0.81	3.45
(p-value)	0.00	0.03	0.00	0.00	0.42	0.00
Session 4	2.44	1.80	3.18	0.21	1.90	2.15
(p-value)	0.01	0.07	0.00	0.84	0.06	0.03
All Sessions	7.51	4.05	7.63	-4.86	0.44	5.65
(p-value)	0.00	0.00	0.00	0.00	0.66	0.00

Note: the table reports results of sign-rank tests of the null hypothesis of equal claims between treatments, based on independent matched observations (within subjects). The number of observations is 24 at session-level and 96 overall.

Table 4: Mean and median expected claims, by treatment and phase

	Phase 1	Phase 2	Phase 3	Phase 4	All Phases
<i>Means</i>					
Treatment 1 (10-10)	182.8	167.5	169.8	167.8	172.0
Treatment 2 (80-80)	131.6	94.8	107.5	107.5	110.4
Treatment 3 (10-80)	136.3	139.3	132.9	135.2	135.9
Treatment 4 (80-10)	134.5	129.6	122.5	99.2	121.4
All treatments	146.3	132.8	133.2	127.4	134.9
<i>Medians</i>					
Treatment 1 (10-10)	200.0	200.0	200.0	200.0	200.0
Treatment 2 (80-80)	95.0	80.0	80.0	80.0	80.0
Treatment 3 (10-80)	130.0	140.0	115.0	87.5	130.0
Treatment 4 (80-10)	100.0	100.0	85.5	80.0	80.0
All treatments	160.0	118.0	100.0	80.5	118.0

Note: Treatment 1: $R = \pm 10$ $\tilde{R} = \pm 10$. Treatment 2: $R = \pm 80$ $\tilde{R} = \pm 80$. Treatment 3: $R = \pm 10$ $\tilde{R} = \pm 80$. Treatment 4: $R = \pm 80$ $\tilde{R} = \pm 10$.

Table 5: Claims and expected claims: cross tabulation

Claims	80	Expected claims		Total
		81-199	200	
<i>All treatments</i>				
80	0.32	0.03	0.03	0.38
81-199	0.08	0.20	0.10	0.37
200	0.03	0.01	0.21	0.25
Total	0.43	0.24	0.34	1.00
<i>Treatment 1 (10-10)</i>				
80	0.04	0.04	0.00	0.08
81-199	0.04	0.22	0.22	0.48
200	0.03	0.03	0.38	0.44
Total	0.11	0.29	0.59	1.00
<i>Treatment 2 (80-80)</i>				
80	0.53	0.03	0.01	0.57
81-199	0.07	0.19	0.04	0.30
200	0.02	0.00	0.10	0.13
Total	0.63	0.22	0.16	1.00
<i>Treatment 3 (10-80)</i>				
80	0.22	0.02	0.02	0.26
81-199	0.13	0.20	0.08	0.41
200	0.07	0.01	0.25	0.33
Total	0.42	0.23	0.35	1.00
<i>Treatment 4 (80-10)</i>				
80	0.49	0.03	0.09	0.61
81-199	0.06	0.18	0.05	0.29
200	0.00	0.00	0.09	0.09
Total	0.55	0.21	0.24	1.00

Note: the table reports observed frequencies for claims (rows) and expected claims (columns). Treatment 1: $R = \pm 10$ $\tilde{R} = \pm 10$. Treatment 2: $R = \pm 80$ $\tilde{R} = \pm 80$.

Treatment 3: $R = \pm 10$ $\tilde{R} = \pm 80$. Treatment 4: $R = \pm 80$ $\tilde{R} = \pm 10$.

Table 6: Claims: tabulations across treatments

	80	81-199	200	
<i>Treatment 10-10</i>				<i>Treatment 80-80</i>
80	0.06	0.02	0.00	0.08
81-199	0.27	0.20	0.01	0.48
200	0.24	0.08	0.11	0.44
Total	0.57	0.30	0.12	1.00
<i>Treatment 10-10</i>				<i>Treatment 10-80</i>
80	0.06	0.01	0.01	0.08
81-199	0.10	0.30	0.07	0.48
200	0.09	0.09	0.25	0.44
Total	0.25	0.40	0.32	1.00
<i>Treatment 10-10</i>				<i>Treatment 80-10</i>
80	0.08	0.00	0.00	0.08
81-199	0.28	0.19	0.01	0.48
200	0.25	0.10	0.08	0.44
Total	0.61	0.29	0.09	1.00

Note: the figures reported are observed frequencies across the 10-10 treatment (rows) and each of the other treatments (columns).

Table 7: Subject types: tabulations across treatments

	NASH	STRA	TEAM	WEAK	MORE	Total
Treatment 10-10						
NASH	0.04	0.00	0.00	0.00	0.00	0.04
STRA	0.00	0.14	0.00	0.00	0.00	0.14
TEAM	0.00	0.00	0.44	0.00	0.00	0.44
WEAK	0.00	0.00	0.00	0.26	0.00	0.26
MORE	0.00	0.00	0.00	0.00	0.13	0.13
Treatment 80-80						
NASH	0.04	0.07	0.21	0.15	0.06	0.53
STRA	0.00	0.02	0.01	0.03	0.00	0.06
TEAM	0.00	0.01	0.11	0.00	0.02	0.15
LESS	0.00	0.00	0.03	0.07	0.01	0.11
MORE	0.00	0.03	0.07	0.01	0.03	0.15
Treatment 10-80						
NASH	0.02	0.02	0.06	0.09	0.02	0.22
STRA	0.00	0.05	0.01	0.00	0.00	0.06
TEAM	0.00	0.02	0.23	0.03	0.02	0.30
LESS	0.00	0.01	0.03	0.09	0.01	0.15
MORE	0.02	0.03	0.10	0.04	0.07	0.27
Treatment 80-10						
NASH	0.04	0.07	0.18	0.11	0.08	0.49
STRA	0.00	0.01	0.00	0.02	0.00	0.03
TEAM	0.00	0.00	0.11	0.02	0.01	0.15
LESS	0.00	0.02	0.09	0.08	0.02	0.22
MORE	0.00	0.03	0.05	0.02	0.01	0.11

Note: the table reports observed frequencies for claims (rows) and expected claims (columns) for the revealed strategy types identified in section 4.3.

Table 8: Differences between treatments and strategy types: OLS

	8080-1010	1080-1010	8010-1010	8080-1080	8080-8010	8010-1080
Overall	-58.9** (-10.84)	-25.3** (-4.96)	-62.1** (-11.47)	-33.5** (-5.87)	3.3 (0.74)	-36.8** (-6.48)
By strategy						
Conditional	-66.1** (-5.25)	-39.4** (-3.18)	-71.2** (-5.40)	-26.8 (-1.90)	5.1 (0.49)	-31.8* (-2.29)
Collusive	-25.2* (-2.06)	-3.9 (-0.32)	-57.7** (-4.51)	-21.4 (-1.56)	32.5** (3.24)	-53.8** (-4.00)
Risk averse	-78.9** (-8.13)	-31.0** (-3.26)	-76.7** (-7.56)	-47.8** (-4.42)	-2.2 (-0.27)	-45.6** (-4.27)
Other	-61.2** (-6.87)	-28.0** (-3.20)	-51.3** (-5.51)	-33.2** (-3.34)	-9.9 (-1.35)	-23.3* (-2.38)
Not available	-18.8 (-0.74)	0.0 (0.00)	-32.3 (-1.22)	-18.8 (-0.67)	13.5 (0.65)	-32.3 (-1.16)
R^2	0.62	0.25	0.60	0.29	0.13	0.34
N. of obs.	96	96	96	96	96	96

Note: Dependent variable: difference within subjects between treatments. * = p<0.05,

** = p<0.01. Heteroskedasticity robust standard errors.

Table 9: Differences between treatments and strategy types: Tobit

	8080-1010	1080-1010	8010-1010	8080-1080	8080-8010	8010-1080
Conditional	-74.8** (-4.60)	-42.3** (-3.03)	-81.1** (-4.91)	-30.5 (-1.92)	5.1 (0.43)	-35.4* (-2.39)
Collusive	-20.9 (-1.87)	-4.1 (-1.14)	-60.3** (-3.52)	-18.7 (-1.40)	35.4** (2.83)	-57.9** (-3.20)
Risk averse	-96.2** (-7.02)	-33.1** (-2.76)	-95.3** (-6.73)	-57.9** (-4.62)	-2.1 (-1.00)	-54.0** (-4.67)
Other	-68.5** (-5.87)	-29.4** (-3.32)	-56.7** (-5.12)	-38.4** (-2.88)	-9.2 (-1.00)	-25.9* (-2.23)
Not available	-18.7 (-1.15)	0.0 (0.00)	-40.0 (-1.18)	-26.0 (-0.67)	13.5 (1.45)	-39.7 (-1.18)
sigma						
Constant	60.2** (14.37)	53.3** (10.58)	64.5** (15.31)	65.2** (10.32)	41.3** (7.95)	62.6** (12.58)
N. obs.	96	96	96	96	96	96

Note: Dependent variable: difference within subjects between treatments. * indicates

p-value <0.05, ** indicates p-value <0.01. Heteroskedasticity robust standard errors.