



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS



University of
Nottingham
UK | CHINA | MALAYSIA

Discussion Paper No. 2019-13

Despoina Alempaki
Andrew M Colman
Felix Kölle
Graham Loomes
Briony D Pulford

November 2019

**Investigating the failure to best
respond in experimental games**

CeDEx Discussion Paper Series

ISSN 1749 - 3293



CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit <http://www.nottingham.ac.uk/cedex> for more information about the Centre or contact

Suzanne Robey
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 14763
suzanne.robey@nottingham.ac.uk

The full list of CeDEX Discussion Papers is available at

<http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx>

Investigating the failure to best respond in experimental games

Despoina Alempaki^{a,*}, Andrew M. Colman^b, Felix Kölle^{c,d,*}, Graham Loomes^a,
Briony D. Pulford^b

This version: November 26th 2019

Abstract: In experimental games, a substantial minority of players often fail to best respond. Using two-person 3x3 one-shot games, we investigated whether ‘structuring’ the pre-decision deliberation process produces greater consistency between individuals’ stated values and beliefs on the one hand and their choice of action on the other. Despite this intervention, only just over half of strategy choices constituted best responses. Allowing for risk aversion made little systematic difference. Distinguishing between players according to their other-regarding preferences made a statistically significant difference, but best response rates increased only marginally. It may be that some irreducible minimum level of noise/imprecision generates some proportion of sub-optimal choices. If so, more research might usefully be directed towards competing models of stochastic strategic choice. (119 words)

Keywords: game theory; best response; strategic thinking; social preferences; stated beliefs

JEL: A13; C72; C91; C92; D84

Acknowledgements: We thank the Leverhulme Trust ‘Value’ programme (RP2012-V-022) and the ESRC’s Network for Integrated Behavioural Science programme (ES/K002201/1 and ES/P008976/1) for financial support, and we thank participants in NIBS workshops and the 12th JDMx in Trento for numerous helpful comments and suggestions.

^a Warwick Business School, University of Warwick, Coventry, CV4 7AL, UK

^b Department of Neuroscience, Psychology and Behaviour, University of Leicester, Leicester, LE1 7RH, UK

^c Centre for Decision Research and Experimental Economics, School of Economics, University of Nottingham, Nottingham, UK

^d Faculty of Management, Economics and Social Sciences, University of Cologne, Cologne, Germany

* Corresponding authors at despoina.alempaki@wbs.ac.uk, koelle@wiso.uni-koeln.de

1. Introduction

The experimental game literature has produced a number of studies showing that a substantial proportion of individuals' strategy choices are not best responses, as judged according to the belief-weighted values of the available options, especially in environments where games are complex and learning opportunities are limited (e.g., see Costa-Gomes and Weizsäcker, 2008; Danz et al., 2012; Hoffman, 2014; Polonio and Coricelli, 2019; Sutter et al., 2013).

Possible reasons include subjects' naivety, low engagement, or limited understanding of the strategic environment, especially when subjects are inexperienced and a game is presented for the first time. Some have suggested that there may be a lack of game-form recognition, i.e., a failure to understand correctly the relationship between possible choices, outcomes and payoffs (e.g., Bosch-Rosa and Meissner, 2019; Cason and Plott, 2014; Chou et al., 2009; Cox and James, 2012; Fehr and Huck, 2016; Rydval et al., 2009; Zonca et al., 2018). Studies using choice process data (e.g., Brocas et al., 2014; Devetag et al., 2016; Hristova and Grinberg, 2005; Polonio et al., 2015; Stewart et al., 2016) suggest that when choosing an action, subjects often pay disproportionately more attention to their own payoffs or to specific salient matrix cells, and a non-negligible fraction of subjects never look at the opponent's payoff, thereby completely disregarding the strategic nature of the game they are playing. As a result, some part of the observed inconsistency might be driven by a failure to incorporate all relevant information, or by the use of heuristic rules that correspond to a simplified - and often incorrect - version of the actual game in question.

Another possible source of the seeming failure to best respond might lie in the existence of other-regarding motives (see e.g., Fehr and Schmidt, 2006; Sobel, 2005 for overviews of the literature). If individuals' choices are not solely driven by self-interest but involve social preferences, it should not be surprising that subjects depart from best responses calculated on the basis of own payoffs only. To date, the role of other-regarding motives in normal form games has been investigated mainly indirectly by monitoring the patterns of information acquisition using eye- or mouse-tracking and connecting the revealed search patterns to different types of other-regarding preferences (e.g., Costa-Gomes et al., 2001; Devetag et al., 2016; Polonio et al., 2015; Polonio and Coricelli, 2019). While evidence from the aforementioned literature suggests that other-regarding motives may interact with the observed levels of strategic sophistication, the correspondence between choice and process data in establishing causal links is less than perfect, since all the inferences are drawn via subjects' information search patterns. Ideally, what is needed is an explicit measure of individuals' other-regarding propensities as related to the payoff structures of the games under consideration.

Another possibility is that many people's preferences are inherently imprecise, so that some degree of variability or 'noise' enters into strategic choice, with the result that in some proportion of cases an option other than the best response may be chosen. Such variability has been widely reported

in the domain of individual choice between pairs of risky monetary lotteries: even though the payoffs of each lottery are familiar sums of money and their respective probabilities are clearly specified and implemented via some simple transparent random mechanism, most subjects have a propensity to choose differently on at least one occasion when the same pair is presented several times in the course of an experimental session (see Bardsley et al., 2010, Chapter 7, for an overview). The implication is that even if an individual's 'best response' to the choice offered to them were to be selected more often than the other option, we would still observe some frequency of non-optimal choices. If that is the case for relatively straightforward binary choices between well-specified lotteries, perhaps there is even greater scope for noise and imprecision to generate non-optimal decisions when the scenarios are more complex: that is, when the outcomes are pairs of own-other payoffs involving some other-regarding preferences, when the decision weights are beliefs about others' behaviour rather than objective explicit probabilities, and when there are three options rather than just two.

The issue we address is whether enabling subjects to become more familiar with the elements of the game they are facing and encouraging them to think more systematically about the subjective values they assign to the payoff pairs and also about the weights they attach to other players' possible actions would increase best response rates. Is it possible to structure the decision environment in such a way that it is easier to make the computations that are assumed to underpin the identification of optimal choices? Would such a manipulation increase best response rates? If so, is this equally true for purely self-interested individuals and for those exhibiting some form of social preferences? Or is it the case that even under such favourable circumstances there will still be some tendency for a proportion of strategy choices to be non-optimal, accounted for, perhaps, by some intrinsic imprecision in people's values and judgments?

In this paper, we investigate these questions using a laboratory experiment in which participants are presented with a set of 8 two-person 3×3 normal-form games. We try to make it as easy as reasonably possible for players to best respond by prompting them to think about their evaluation of payoffs in conjunction with their beliefs about the other players' possible choices, after which they choose their own strategies in the light of those deliberations and with that information still in front of them.

Specifically, using a *ranking* task, we first ask subjects to consider the subjective value they attach to the various cells in each 3×3 game. Based on subjects' stated rankings, we can infer something about their other-regarding motives, thus relaxing the assumption of pure self-interest (as used by most previous literature) in cases where it does not seem appropriate. Some previous studies (e.g., Bayer and Renou, 2016) have tried to infer subjects' social preferences from their behaviour in a scenario such as a modified dictator game and then 'import' this information into the games that are their main focus. However, other-regarding motives are highly context-dependent and findings from a different decision environment might not carry over (see Galizzi and Navarro-Martinez, 2019, for evidence and a meta-

analysis). By contrast, our ranking task explores the expression of other-regarding preferences in a format which is *intrinsic* to the strategic situation faced in the actual games under consideration.

Second, using a *belief* task, we ask subjects to quantify the chances of each of the opponents' strategies being played. Importantly, we do not require those estimates to conform to any particular assumptions about the rationality of reasoning of the other players; we simply ask subjects for their best judgements. Previous literature investigating the effect of belief elicitation on equilibrium play has obtained mixed results with inconclusive evidence: some studies find that belief elicitation does not affect game play (e.g., Costa-Gomes and Weizsäcker, 2008; Inavov, 2011; Nyarko and Schotter, 2002; Polonio and Coricelli, 2019), whereas others (e.g., Hoffmann, 2014) show that whether belief elicitation influences game play depends on the properties of the game (see Schlag et al., 2015 and Schotter and Trevino, 2014, for reviews). In contrast to these studies, here we investigate the effect of beliefs in conjunction with stated rankings over payoffs, rather than the effect of beliefs alone.

To see how far and in what ways our intervention alters behaviour, we compare the patterns of strategy choice in what we shall call the *Structured* sample with the responses of a different sample drawn from the same population who we shall call the *Unstructured* sample and who were presented with exactly the same games but were asked to make their decisions without any prior structured tasks. Since the great majority of game experiments to date have been conducted in this unstructured manner, it is of interest to see how, if at all, the patterns of choice might be affected.

Our main results can be summarized as follows. In line with previous literature (Costa-Gomes and Weizsäcker, 2008; Danz et al., 2012; Hoffman, 2014; Polonio and Coricelli, 2019; Sutter et al., 2013) we find that a sizeable minority of players fail to best respond to their own stated beliefs. While we find that the level of consistency increases significantly when we make some allowance for other-regarding preferences as revealed by the ranking task, the observed difference is relatively small (54% vs. 57%). In fact, if we consider only the subset of individuals that exhibit no other-regarding concerns in the ranking task, their best response rate is unaffected (55%). For subjects we classify as inequity averse, in contrast, best response rates increase significantly from 54% to 61% when using rankings instead of own payoffs. This shows that the observed increase at the aggregate level is almost entirely driven by those subjects who are not motivated only by their own earnings. This is reassuring, as it shows that the ranking task is picking up something that feeds into subjects' strategy choices.

Our results also suggest that the likelihood of choosing non-optimally is decreasing in the costs of doing so. That is, while non-optimal choices are relatively common when the expected payoffs from the different strategies are very similar, such choices become less and less likely as the difference in expected payoffs between strategies grows, whether measured in terms of foregone monetary payoffs or in terms of foregone ranking values. This appears to be compatible with the notion of Quantal Response Equilibrium (McKelvey and Palfrey, 1995) or some other 'error' model (see Harrison, 1989,

for a discussion of the “flat maximum” critique). Finally, when comparing patterns of choice in the Structured and Unstructured treatments, we find overall no impact of structuring the decision process: our attempts to induce “harder thinking” did not cause subjects to choose more sophisticated strategies.

The rest of the paper is organized as follows. In Section 2, we describe the design and implementation of the experiment. We present our main results in Section 3. Section 4 discusses and concludes.

2. The experiment

We chose a set of eight 3×3 normal-form games that were adjusted versions of the games used in Colman et al. (2014). The games are displayed in Figure 1.¹ In each cell, the first number indicates the payoff to the BLUE (row) player and the second number indicates the payoff to the RED (column) player. The payoffs in each cell were sums of money (in UK pounds). The games were chosen because they have no obvious payoff-dominant solutions and because they were explicitly designed to differentiate between competing theoretical explanations (see Table 1). Furthermore, previous evidence from the patterns of strategy choices in these games suggests relatively low effort thinking, which gives us enough room to detect whether structuring responses leads to higher levels of reasoning (see Colman et al., 2014; Pulford et al., 2018).

Table 1 summarizes the strategic structure for each game and player role. Besides the Nash equilibrium prediction, we consider additionally Level- k models, which often out-predict equilibrium play (e.g., Camerer et al., 2004; Costa-Gomes and Crawford, 2006; Ho et al., 1998; Nagel, 1995; Stahl and Wilson, 1994, 1995) to allow for bounded depth of reasoning. The *Level-1* model assumes that a player best responds to the belief that assigns uniform probabilities to their counterparts’ actions. The *Level-2* model predicts that a player best responds to the belief that their counterpart is playing according to the Level-1 model.

At the beginning of the experiment, subjects were randomly allocated either the role of the BLUE (row) or the RED (column) player, and they remained in that role throughout the whole experiment. Participants were then presented with each game in turn, proceeding at their own speed. The order in which the eight games were displayed was randomized and subjects received no feedback about others’ choices until the end of the experiment.

¹ Our games differ from Colman et al. (2014) as follows. We doubled all the original payoffs in order to bring earnings more into line with other studies in this literature (relatedly Pulford et al., 2018 multiply all payoffs by five and find no evidence of a stake size effect). We also substituted the original Game 8 with a new game, because the original Game 8 yielded similar predictions to Game 6.

Figure 1. Experimental games. Underlined payoffs indicate the Nash equilibria in pure strategies.

1	D	E	F	2	D	E	F
A	8, 6	2, 8	2, 0	A	<u>6, 6</u>	2, 0	0, 4
B	10, 0	6, 6	0, 8	B	2, 2	<u>4, 4</u>	6, 0
C	2, 2	8, 0	<u>4, 4</u>	C	4, 0	2, 4	4, 10

3	D	E	F	4	D	E	F
A	6, 6	0, 10	2, 0	A	<u>6, 6</u>	2, 2	0, 4
B	10, 0	4, 4	0, 8	B	2, 2	2, 8	6, 0
C	0, 2	8, 0	<u>4, 4</u>	C	0, 0	4, 2	4, 10

5	D	E	F	6	D	E	F
A	4, 10	0, 2	6, 2	A	8, 6	4, 8	10, 4
B	0, 4	<u>6, 6</u>	0, 4	B	0, 0	8, 4	6, 8
C	8, 2	4, 2	0, 6	C	<u>10, 10</u>	2, 8	4, 2

7	D	E	F	8	D	E	F
A	2, 8	0, 2	10, 4	A	<u>6, 6</u>	0, 4	8, 4
B	0, 2	<u>6, 6</u>	0, 2	B	4, 6	10, 4	0, 0
C	6, 4	4, 2	4, 6	C	4, 2	8, 8	4, 10

Table 1. Structure of the games and models' predicted actions.

Game	Unique Nash	Nash Pareto dominated	Symmetric	Nash	Level-1	Level-2
1	✓	✓	×	C-F	B-E	C-F
2	×	×	×	A-D	B-F	B-E
3	✓	✓	✓	C-F	B-E	C-F
4	✓	×	×	A-D	B-F	B-E
5	✓	×	×	B-E	C-D	C-F
6	✓	×	×	C-D	A-E	B-E
7	✓	×	×	B-E	C-D	C-F
8	✓	✓	×	A-D	C-E	B-F

Notes. We indicate whether a game has a unique Nash equilibrium or not (Unique Nash), whether another cell constitutes a strict Pareto improvement (Nash Pareto dominated), whether it is symmetric or not (Symmetric), as well as predictions according to the Nash, Level-1 and Level-2 models. For Game 2 that has multiple Nash equilibria we report predictions on the Pareto optimal.

For each game, subjects completed three different tasks: a *ranking*, a *belief* and a *choice* task.² The purpose of the first two tasks was to structure subjects' decision making process and to induce them to think about all relevant aspects of a game before choosing their preferred strategy. In all tasks, we recorded how much time subjects spent before submitting their decisions. To make each response incentive compatible, it was explained that one out of the eight games would be randomly selected for payment, followed by another random draw to determine whether subjects' earnings were determined according to the ranking, the belief, or the choice task. We now describe each of those tasks in turn, together with the mechanism designed to motivate subjects to give thoughtful and accurate responses.

In the *ranking* task, subjects were asked to rank all possible unique payoff combinations in a particular game from their most preferred to their least preferred one. Figure 2 shows a screenshot of the interface of this task as shown to the subjects.

Figure 2. Screenshot of the ranking task

You are the **BLUE** decision maker. Please make your Type I decision.

Grid #1	D	E	F
A	8, 6	4, 8	10, 4
B	0, 0	8, 4	6, 8
C	10, 10	2, 8	4, 2

Type I decision	Rank from 1 to 9
8, 6	
4, 8	
10, 4	
0, 0	
8, 4	
6, 8	
10, 10	
2, 8	
4, 2	

OK

At the top of the screen, subjects were shown the game being played. In the bottom left corner of the screen, they were shown all the possible payoff pairs, ordered as they appear in the game from top left to bottom right.³ They stated their ranking of these pairs by typing in a number between 1 and 9, where 1 corresponds to the pair they liked best and 9 indicated their least preferred pair. Subjects had to state a complete and strict monotonic ranking: i.e., they were not allowed to state indifference.

² In the experiment, the different tasks were called Type I, Type II, and Type III decisions respectively (see Figure 2 – 4, as well as the experimental instructions in Appendix C).

³ The order of the payoffs from top left to bottom right was per row for the BLUE player and per column for the RED player in order to correspond to their actual strategies in the *choice* task.

In the event that the ranking task was selected for payment, one player in a pair of players (either RED or BLUE) was selected as the decision maker. We then randomly picked two of the possible payoff combinations from the selected game and paid both players according to the combination that the randomly selected decision maker had ranked as more preferable.

Once subjects had submitted their ranking, they proceeded automatically to the *belief* task. In this task, players were asked to think about the ten players participating in the same session who had been assigned to the role of the other colour and they were asked to state their best estimates about how many of these ten players would choose each of their three possible strategies available to them.⁴ A screenshot of the interface used in the belief task is shown in Figure 3.

Figure 3. Screenshot of the belief task

You are the **BLUE** decision maker. Please make your Type II decision.

Grid #1	D	E	F
A	8, 6	4, 8	10, 4
B	0, 0	8, 4	6, 8
C	10, 10	2, 8	4, 2

Type I decision	
Option	Rank from 1 to 9
8, 6	3
4, 8	5
10, 4	2
0, 0	9
8, 4	4
6, 8	6
10, 10	1
2, 8	8
4, 2	7

Type II decision
My best estimate of how many of the 10 RED decision makers will choose D is <input type="text"/>
My best estimate of how many of the 10 RED decision makers will choose E is <input type="text"/>
My best estimate of how many of the 10 RED decision makers will choose F is <input type="text"/>

In the event that the belief task was selected for payment, we randomly picked one of the three strategies available to the other colour of player and then compared the subject's stated belief about the frequency of that choice with the actual frequency among the ten players assigned to that colour in the session. If both numbers matched, subjects were paid £5. Otherwise they received no payoff.⁵

Finally, on the last screen of each game subjects had to indicate their preferred strategy in the *choice* task (see Figure 4 for a screenshot of the interface) on the usual understanding that if this task were selected as the basis for payment, they would be paired at random with another participant and each member of the pair would be paid according to the intersection of their strategy choice.

⁴ All Structured sessions were conducted with 20 participants each, 10 BLUE and 10 RED players.

⁵ We chose this incentive mechanism instead of the quadratic scoring rule because of its simplicity and to avoid distortion due to the possibility of participants reporting beliefs away from extreme probabilities (see the discussion in Schlag et al., 2015).

Figure 4. Screenshot of the choice task

You are the **BLUE** decision maker. Please make your Type III decision.

Grid #1	D	E	F
A	8, 6	4, 8	10, 4
B	0, 0	8, 4	6, 8
C	10, 10	2, 8	4, 2

Type I decision

Option	Rank from 1 to 9
8, 6	3
4, 8	5
10, 4	2
0, 0	9
8, 4	4
6, 8	6
10, 10	1
2, 8	8
4, 2	7

Type II decision

My best estimate of how many of the 10 **RED** decision makers will choose **D** is 2

My best estimate of how many of the 10 **RED** decision makers will choose **E** is 3

My best estimate of how many of the 10 **RED** decision makers will choose **F** is 5

Type III decision

I choose ☐ A ☐ B ☐ C

OK

Importantly, in order to reinforce the effect of previous deliberation on the selected strategy and to control for differences in working memory capacity (see e.g. Devetag et al., 2016), at the time they were choosing their strategy, subjects could see their responses in the ranking and in the belief tasks (as shown in Figure 4), making it as easy as we could for them to weight the values they placed on payoffs by their beliefs about the other players' actions, if that is what they wished to do.

In order to judge whether our Structured manipulation had any substantial systematic effect on strategy choice, we ran a separate Unstructured control treatment: using the exact same games, subjects were simply asked to state – without any prior deliberation tasks – which of the three available strategies they wanted to play. In this treatment, we simply picked one game at random and then paid subjects according to the intersection of their choices.⁶

The experiment was run at the CeDEx laboratory at the University of Nottingham using students from a wide range of disciplines recruited through ORSEE (Greiner, 2015). The experiment was computerized using z-Tree (Fischbacher, 2007). We conducted ten sessions (five per treatment) with a total of 194 subjects (100 in the Structured treatment, 94 in the Unstructured treatment, 61% of them female, average age 20.8 years). At the beginning of each session, subjects were informed about their role (BLUE or RED) and about the payment procedure that would follow at the end of the experiment. Before the experiment started, participants were asked to read some preliminary instructions of an example 3x3 game and to demonstrate they fully understand the required tasks by answering a series of

⁶ In this Unstructured treatment, the sequence of the eight games was repeated twelve times. In the present paper we discuss only the data from the first sequence. Since subjects had no information about the number of sequences, the repetition could have not affected their choices when they saw each game for the first time. Further details about the patterns of behaviour in subsequent repetitions is available from the corresponding authors on request.

control questions before they could proceed to the actual experiment. Sessions lasted approximately an hour, and average earnings were about £6.90 (including a £3 show-up fee).

3. Results

We organise our results as follows. In Section 3.1, we start with a descriptive analysis of choices and stated beliefs in the Structured treatment and discuss to what extent subjects best respond to their beliefs assuming they are only interested in maximizing their own payoff. In Section 3.2, we then turn to the ranking task, analysing the extent to which subjects are motivated by other-regarding concerns. We then investigate whether accounting for such social preferences increases the proportion of optimal choices. In Section 3.3, we discuss possible determinants of non-optimal behaviour. Finally, to test whether structuring subjects' decision process had any effect on behaviour, in Section 3.4 we compare patterns of chosen strategies in the choice task across the Structured and the Unstructured treatments.

3.1 Beliefs, choices and best response

The aggregate data in Table 2 show that the fraction of Nash equilibrium choices varies considerably across games with a minimum of 0.09 (Game 7) and a maximum of 0.49 (Game 6). The average rate of equilibrium choices in games with a unique Nash equilibrium (all but Game 2) is equal to 0.27, which is significantly lower than would be predicted by chance (t-test, $p = 0.003$). Table 2 further reveals that, in line with previous evidence (e.g., Costa-Gomes and Weizsäcker, 2008; Polonio and Coricelli, 2019), the Level-1 and Level-2 models both outperform equilibrium predictions. The fractions of choices consistent with the Level-1 and Level-2 prediction amount to 0.50 and 0.41, respectively, on average, compared with 0.27 for the Nash equilibrium.

At the aggregate level, a similar pattern emerges for beliefs. Most participants do not expect their counterparts to play according to equilibrium: the average probability with which participants estimated that the counterpart is playing the equilibrium strategy (in games with a unique Nash equilibrium) is equal to 0.28. Instead, it is most often predicted that the counterpart would choose according to the Level-1 model (probability = 0.52), followed by the Level-2 model (probability = 0.34).

The fact that both choices and beliefs are best described by the Level-1 model already provides first aggregate evidence that subjects do not always best-respond. In particular, the proportion of Level-2 choices is much lower than the proportion of Level-1 beliefs (see Polonio and Coricelli, 2019 for similar evidence). To provide more conclusive evidence on subjects' best response behaviour, in the following, we investigate the level of consistency between choices and stated beliefs at the individual level. To test whether subjects best respond to their stated beliefs, we calculate a player's expected payoff for each possible strategy available to them on the basis of their stated beliefs, assuming either linear utility of payoffs, or some degree of risk aversion. More specifically, we use the power law

function x^α with $\alpha = 1$, $\alpha = 0.8$, and $\alpha = 0.5$. We then simply count how often a subject chooses the strategy that gives them the highest expected utility. The results are given in Table 3.

Table 2. Consistency of beliefs and choices with theoretical predictions

Game	Choice task			Belief task		
	Nash	Level-1	Level-2	Nash	Level-1	Level-2
1	0.31	0.53	0.31	0.32	0.46	0.32
2	0.33	0.44	0.46	0.26	0.56	0.38
3	0.42	0.42	0.42	0.26	0.52	0.26
4	0.26	0.28	0.59	0.29	0.43	0.48
5	0.17	0.63	0.50	0.14	0.64	0.40
6	0.49	0.45	0.27	0.51	0.40	0.20
7	0.09	0.62	0.58	0.12	0.63	0.44
8	0.15	0.67	0.18	0.29	0.48	0.22
Average	0.27	0.50	0.41	0.28	0.52	0.34

Notes: The left side shows the average proportion of choices in accordance with the different models' predictions. The right side shows the stated beliefs (average probabilities on model predictions) about the model of choice of the counterpart. For Game 2, which has two Nash Equilibria, we show the rate with which subjects chose the Pareto dominant Nash Equilibrium. Tables A1 and A2 in Appendix A report these rates separately for the row and column players. The total average for the Nash equilibrium is calculated over the games with unique equilibrium (excluding Game 2).

Table 3 reveals that with linear utilities of payoffs ($\alpha = 1$), the average proportion of best responses varies from a minimum of 0.42 (Game 8) to a maximum of 0.68 (Game 7). Averaged over all games, the best response rate is 0.54. Although this is significantly higher than predicted by chance (t-test, $p < 0.001$), it means that almost half of all strategy choices are non-optimal. Furthermore, the average best response rate does not change if we allow for some degree of risk aversion (see last two columns of Table 3). These data are in line with results from previous literature (e.g., Costa-Gomes and Weizsäcker, 2008; Polonio and Coricelli, 2019; Rey-Biel, 2009; Sutter et al., 2013), which report consistency levels ranging from 54% to 67%. That is, despite having their personal rankings of payoff pairs and their stated beliefs about their counterparts' strategies on display at the time they are choosing their own strategy, participants often choose options which do not maximise the expected utility of their own payoffs.

Table 3. Frequency of best responses using expected payoffs

Game	$U(x) = x$	$U(x) = x^{0.8}$	$U(x) = x^{0.5}$
1	45%	47%	46%
2	54%	52%	55%
3	49%	51%	50%
4	53%	52%	50%
5	55%	53%	51%
6	67%	64%	63%
7	68%	65%	64%
8	42%	46%	54%
Average	54%	54%	54%

3.2 The role of other-regarding preferences

A possible explanation for the seeming failure to best respond as described above might be that individuals do not only care about their own payoff, but also incorporate the payoffs to others into their utility function (see, for example, Sobel, 2005, for an overview of the literature at that time). It could be that choices which appear non-optimal under the assumption of pure self-interest might be fully rational once we allow for subjects' other-regarding preferences.

To examine this possibility, we turn to the results of the ranking task, in which subjects in the Structured treatment were asked to rank the different payoff combinations in each game from most preferred (1) to least preferred (9). Table 4 shows the mean ranks for all own-other payoff pairs, averaged over all games. Not surprisingly, subjects generally prefer more money over less. That is, holding constant the other player's payoff, the mean ranking score generally decreases as own payoff increases.

The results also reveal that subjects on average are inequity averse (Bolton and Ockenfels, 2000; Fehr and Schmidt, 1999). That is, in Table 4, holding constant the subject's own payoff (i.e., fixing a row), the most preferred pair lies on the main diagonal (as highlighted by the grey shaded cells) where both players obtain the same positive payoff. The exception occurs in the first row where both payoffs are zero: on average, when their own payoff is zero, subjects prefer unequal payoffs even though this involves the other player receiving more than they do.

Table 4. Mean ranking for pairs of own and other's payoffs

Own payoff	Other's payoff					
	0	2	4	6	8	10
0	8.4	7.3	7.8	7.2	7.5	7.0
2	6.5	5.7	6.5	6.0	6.4	-
4	5.2	5.0	3.8	4.9	6.2	4.9
6	2.9	3.0	3.3	2.1	4.3	-
8	3.1	2.3	3.5	2.4	1.7	-
10	1.9	-	1.7	-	-	1.1

Notes: Ranking from 1 (best) to 9 (worst). Payoff pairs that did not appear in any of our games are displayed by “-”. In some games there were less than nine unique payoff pairs as some payoff pairs appeared multiple times. In particular, in games 3, 4, and 5 there are only eight unique payoff pairs, and in game 7 there are only seven. In this case, subjects had to rank each payoff pair only once and hence, only ranks between 1 and 8 or 1 to 7, respectively, were possible. To correct for this, in Table A3 in Appendix A we display an alternative version of Table 3 in which we normalize the rankings. The results remain effectively unchanged.

These results are further corroborated by OLS regressions, in which we use the rank as the dependent variable and own payoff as well as the absolute difference between own and other's payoff as independent variables. The results are reported in Table 5. They confirm that increasing own payoff has a strong and significant negative effect on stated ranks, consistent with people preferring more money over less. At the same time, the absolute difference between their own and their counterpart's payoff has a significant positive effect, indicating that, *ceteris paribus*, people dislike inequality.

Table 5. OLS regressions: Determinants of ranking

Dependent variable: Ranking	
Own payoff π_i	-0.652*** (0.010)
Absolute payoff difference $ \pi_i - \pi_j $	0.106*** (0.012)
Constant	7.127*** (0.072)
# Observations (clusters)	6700 (100)
R ²	0.689

Notes: This table reports coefficient estimates from an OLS regression. Standard errors clustered at the individual level are reported in parentheses. *** p < 0.01

To explore whether such other-regarding concerns might account for some or many of the departures from own-payoff best responses, we re-calculate optimal choices based on subjects' stated beliefs and rankings (rather than payoffs). That is, similar to the analysis above, we first calculate the

expected ranking for each possible strategy, assuming linear utilities in rankings. We then simply count how often a subject chooses the strategy that gives her the most preferred expected ranking.

The results reveal that best-response rates do indeed increase relative to the case when only own payoffs are considered. In particular, in 6 out of 8 games the fraction of best response rates is higher when using subjects' rankings rather than their own payoffs, with the difference between the two ranging from one to eleven percentage points (compare Table A4 in the appendix). On average, however, the best response rate increases only moderately from 54% to 57%, a difference that is nevertheless statistically significant (Signrank test, $p = 0.004$; paired t-test, $p = 0.025$).⁷

To shed some further light on the role of other-regarding motives, we explore the underlying heterogeneity in social preferences. In particular, while the analysis above indicates that subjects are on average inequity averse, previous literature has shown that individuals typically differ with regard to their other-regarding concerns (see e.g., the discussion in Iriberri and Rey-Biel, 2013). To test for this, as a simple measure of a subject's social type, we re-estimate the model from Table 5 separately for each individual. We then use the sign and the significance of the coefficient for the absolute difference between own and other's payoffs to classify subjects into different distributional preference types: *Selfish* (if a subject's ranking is not significantly affected by differences between own and other's payoffs), *Inequity Averse* (if a subject's ranking is significantly increasing in payoff differences), and *Inequity Seeking* (if a subject's ranking is significantly decreasing in payoff differences).

On this basis, 55 out of 100 subjects are classified as Selfish while 44 subjects are classified as Inequity Averse.⁸ We find no subject to be Inequity Seeking.⁹ This classification allows us to recalculate the best response rates separately for each type. For Selfish individuals we find that, on average, the best response rate amounts to 55%, irrespective of whether using expected payoffs or expected rankings (Signrank test, $p = 0.372$; paired t-test, $p = 0.766$). For individuals classified as Inequity Averse, in contrast, we find that including their responses in the ranking task significantly increases their level of best response from 54% to 61% (Signrank test, $p = 0.005$; paired t-test, $p = 0.011$). This shows that the observed increase in best response rates at the aggregate level when using rankings instead of own payoffs is almost entirely driven by those subjects who are not motivated only by their own earnings. This is reassuring as it shows that the ranking task is picking up something that feeds into subjects' strategy choices.

⁷ For the test we use an individual's average over all games as the unit of observation.

⁸ One subject falls under neither of these categories as he/she did not significantly react to a change in own payoffs. In the following, we discard this subject from our analysis. However, all our results are robust to the inclusion of this subject into either of the two categories.


⁹ As a robustness check, we also applied different classification procedures. In particular, we conducted a similar analysis as above, but, following the model of Fehr and Schmidt (1999), allowed for differences in the degree of advantageous and disadvantageous inequity aversion. The results are very similar and available upon request.

In sum, while we find that taking into account individuals' social preferences can indeed improve predicted game play, the differences are only small. Our data therefore give rather limited support for the idea that social preferences are a major explanation for non-optimal behaviour, as conventionally judged in terms of maximising expected own-payoff values.

3.3 Possible factors associated with non-optimal play

In this section, we try to shed some light on possible determinants of non-optimal play other than social preferences. As a first step, we provide some descriptive statistics of the underlying heterogeneity of non-optimal play. At the individual level, we find substantial variation in the degree to which subjects best respond. While the majority of people (67% when using expected payoffs and 77% when using expected rankings) choose optimally in at least half of the games, only very few people choose optimally in all eight games (see Figure A1 in the appendix for the full distribution). The mean (median) number of optimal choices is 4.33 (4) when using expected payoffs, and 4.59 (4.5) when using expected rankings, a difference that is statistically significant (Signrank test, $p = 0.004$; paired t-test, $p = 0.025$). This confirms that taking into account subjects' social preferences increases best response behaviour, but that this effect is only moderate in size.

Figure 5. Percentage of non-optimal choices as a function of foregone expected payoffs (left panel) and foregone expected rankings (right panel)



Next, we look at the cost of non-optimal play. We compare the expected payoff (expected ranking) between the chosen option and the option that would have been optimal given a subject's stated

beliefs. Figure 5 shows the percentage of non-optimal choices as a function of the foregone expected payoff (left panel) or foregone expected ranking (right panel). It appears that non-optimal strategies are particularly likely when the loss resulting from such choices is small, while they become less and less likely the larger the size of the loss. On average, conditional on choosing non-optimally, subjects forego 2 pounds of expected payoffs (median 1.8) and 1.4 points in expected rankings (median 1).

Table 6. Regression analysis of optimal choices

Dependent variable:	Optimal Choice Payoffs			Optimal Choice Rankings		
	(1)	(2)	(3)	(4)	(5)	(6)
Std. dev. in expected payoffs across options	0.201*** (0.070)	0.286*** (0.077)	0.221*** (0.076)			
Std. dev. in payoffs within optimal choice		-0.359*** (0.090)				
Optimal choice contains lowest payoff (1 if yes, 0 otherwise)			-0.896*** (0.191)			
Optimal choice contains highest payoff (1 if yes, 0 otherwise)			0.090 (0.165)			
Std. dev. in expected rankings				0.210** (0.095)	0.275*** (0.101)	0.213*** (0.106)
Std. dev. in ranks within optimal choice					-0.082** (0.032)	
Optimal choice contains least preferred option (1 if yes, 0 otherwise)						-0.687*** (0.190)
Optimal choice contains most preferred option (1 if yes, 0 otherwise)						0.081 (0.214)
Constant	-0.310** (0.151)	0.522** (0.252)	0.194** (0.197)	-0.106 (0.150)	0.207 (0.190)	0.077 (0.162)
# Observations (clusters)	761 (100)	761 (100)	761 (100)	767 (100)	767 (100)	767 (100)
(Pseudo) R ²	0.009	0.027	0.039	0.006	0.015	0.024

Notes: This table reports coefficient estimates from logistic regressions. The dependent variable is whether the choice is optimal regarding the expected payoff (Models (1) – (3)) and the expected ranking (Models (4) – (6)). We only use data from cases in which the optimal choice was unique, i.e., we are excluding cases in which based on subjects' stated beliefs two or more options were optimal. When using expected payoffs this is leaving us with 761 out of 800 cases. When using expected rankings this is leaving us with 767 out of 800 cases. Standard errors clustered at the individual level are reported in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.10.

To provide more detail, Table 6 reports results from a series of logistic regressions with choosing optimally as the binary dependent variable. In Model (1), we use the standard deviation in the expected earnings across the three available strategies as the explanatory variable. The results show a significant positive coefficient, indicating that the more dissimilar the available strategies are (with regard to their expected earnings), the higher the likelihood of choosing optimally: intuitively, if one strategy stands out as the best, the easier it is to choose optimally. As we show in the appendix, this is

further reflected in response times: the bigger the advantage of the best option, the faster subjects reach a decision (see Table A5 in Appendix A).

In Model (2), we add the standard deviation of the three possible own earnings within the optimal choice as a second explanatory variable. The coefficient is significantly negative, indicating that the higher the variation in own payoffs within the choice that is optimal given the stated beliefs, the less likely subjects are to choose optimally. Our finding that strategy variance might act as a determinant of choice is in line with previous studies (see e.g., Devetag et al., 2016; Guida and Devetag, 2013) showing that choice behaviour is susceptible to the influence of out-of-equilibrium features of the games under consideration. Guida and Devetag (2013), for example, show that increasing the payoff variance in the strategy with the highest expected payoff significantly shifts choice behaviour away from that strategy. It is not clear, however, from the aforementioned studies, whether this shift reflects a tendency to pick a strategy that is both attractive and relatively safe or whether it is simply an attempt to avoid the worst possible payoff.

In model (3), we separate this effect by including two dummy variables indicating whether the optimal choice contains the lowest or highest possible payoff within a given game. The results reveal that containing the minimum possible payoff has a strong negative impact on the likelihood of choosing optimally. At the same time, containing the maximum possible payoff only has a small positive and insignificant effect. It thus seems that the negative effect of variation in own payoffs is mainly driven by subjects trying to avoid the worst possible payoff, even if this means deviating from the optimal strategy. In Models (4) to (6), we repeat the same analysis but now using optimal choices calculated based on subjects' rankings of payoffs rather than their own payoffs only. The results corroborate our previous findings.

3.4 Structuring the decision process has no significant effect on strategy choices

Besides trying to understand what determines subjects' best response behaviour, we were further interested in whether structuring subjects' decision-making process by helping them to "think harder" about the game at hand, has any influence on the chosen strategies and their depth of reasoning. To test this, we compare the patterns of choice in the Structured treatment with those in the Unstructured treatment.

Overall, we find very little evidence that structuring players' decision processes has a significant effect on actual game play. That is, we find no significant differences in the rate with which participants play according to the Nash, Level-1, or Level-2 predictions across the two treatments. On average, choices in the Unstructured treatment are actually slightly more likely to be consistent with the Nash prediction in games with a unique Nash equilibrium (31% vs. 27%) and slightly less likely to be consistent with the Level-1 (46% vs. 50%) and Level-2 (38% vs. 41%) prediction, but none of these differences is statistically significant (all p-values > 0.139, compare Table B1 in Appendix B). These

results hold if we compare choices separately for each game and player type (row or column player) (see Tables B2 and B3 in Appendix B). In sum, in line with the results of Costa-Gomes and Weizsäcker (2008) who find no effect of belief elicitation on actual game play, we find no strong effect even when we elicit payoff rankings in conjunction with beliefs.

4. Concluding remarks

The notion of best response is central to the analysis of strategic behaviour, both in conventional game theory and in cognitive hierarchy / Level-k models which allow for more limited depths of reasoning. So it has troubled experimental researchers that numerous studies have reported substantial evidence of failures to best respond.

Our study sought to examine the extent to which best response rates might increase if subjects' possible unfamiliarity with the strategic environment were offset by asking them to focus in turn on their ranking of payoff pairs and on their beliefs about the other players' probable choices before selecting their strategies. Structuring the decision process in this way did not increase best response rates when judged in terms of maximising own-payoff expected values. Nor were rates changed much by allowing for different levels of risk aversion using a standard utility function form (although avoiding the worst possible payoff did appear to carry some weight). We further examined the possible role of subjects' other-regarding considerations: we found that making allowance for such preferences increased best response rates, but only marginally.

It is hard to imagine what more we might have done, within the usual experimental time and budget constraints, to have made it even easier for participants to best respond. The fact that, even given such favourable conditions, just under half of the chosen strategies were non-optimal might be taken to indicate some intrinsic limit to the precision with which the expected utilities of options can be judged by players. The evidence that the likelihood of non-optimal responses falls as the opportunity loss increases is consistent with various models of noise and stochastic behaviour. Quantal Response Equilibrium is probably the best-known of these, but others that adapt random preferences to games (Bardsley et al., 2010, section 7.3.2) or that apply accumulator mechanisms (Golman et al., in press) may also be candidates for further consideration. The persistence of stochastic behaviour even under our Structured condition may provide further impetus to develop and incorporate appropriate 'error' specifications into strategic decision modelling.

Finally, the lack of any significant difference in the overall patterns of choice between Structured and Unstructured treatments may reassure researchers that simply asking participants to make their choices is an adequate way for experiments to be conducted. Dispensing with demanding prior ranking and belief elicitation procedures does not, on this evidence, greatly affect the quality of the data. Scarce laboratory time and money can instead be devoted to collecting larger and more

powerful datasets: not least, ones that would enable us to investigate further the stochastic component in strategic choice.

References

- Bardsley, N., Cubitt, R., Loomes, G., Moffatt, P., Starmer, C. & Sugden, R. (2010). *Experimental economics: rethinking the rules*. Princeton University Press.
- Bayer, R. C., & Renou, L. (2016). Logical abilities and behavior in strategic-form games. *Journal of Economic Psychology*, 56, 39-59.
- Bolton, G. E., & Ockenfels, A. (2000). ERC: A theory of equity, reciprocity, and competition. *American Economic Review*, 90(1), 166-193.
- Bosch-Rosa, C., & Meissner, T. (2019). The One Player Guessing Game: A diagnosis on the relationship between equilibrium play, beliefs, and best responses. Working paper.
- Brocas, I., Carrillo, J. D., Wang, S. W., & Camerer, C. F. (2014). Imperfect choice or imperfect attention? Understanding strategic thinking in private information games. *Review of Economic Studies*, 81(3), 944-970.
- Camerer, C. F., Ho, T. H., & Chong, J. K. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3), 861-898.
- Cason, T. N., & Plott, C. R. (2014). Misconceptions and game form recognition: Challenges to theories of revealed preference and framing. *Journal of Political Economy*, 122(6), 1235-1270.
- Chou, E., McConnell, M., Nagel, R., & Plott, C. R. (2009). The control of game form recognition in experiments: Understanding dominant strategy failures in a simple two person “guessing” game. *Experimental Economics*, 12(2), 159-179.
- Cox, J. C., & James, D. (2012). Clocks and trees: Isomorphic Dutch auctions and centipede games. *Econometrica*, 80(2), 883-903.
- Colman, A. M., Pulford, B. D., & Lawrence, C. L. (2014). Explaining strategic coordination: Cognitive hierarchy theory, strong Stackelberg reasoning, and team reasoning. *Decision*, 1(1), 35.
- Costa-Gomes, M. A., & Crawford, V. P. (2006). Cognition and behavior in two-person guessing games: An experimental study. *American Economic Review*, 96(5), 1737-1768.
- Costa-Gomes, M., Crawford, V. P., & Broseta, B. (2001). Cognition and behavior in normal-form games: An experimental study. *Econometrica*, 69(5), 1193-1235.
- Costa-Gomes, M. A., & Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *The Review of Economic Studies*, 75(3), 729-762.
- Danz, D. N., Fehr, D., & Kübler, D. (2012). Information and beliefs in a repeated normal-form game. *Experimental Economics*, 15(4), 622-640.
- Devetag, G., Di Guida, S., & Polonio, L. (2016). An eye-tracking study of feature-based choice in one-shot games. *Experimental Economics*, 19(1), 177-201.
- Fehr, D., & Huck, S. (2016). Who knows it is a game? On strategic awareness and cognitive ability. *Experimental Economics*, 19(4), 713-726.
- Fehr, E., & Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3), 817-868.
- Fehr, E., & Schmidt, K. M. (2006). The economics of fairness, reciprocity and altruism—experimental evidence and new theories. *Handbook of the economics of giving, altruism and reciprocity*, 1, 615-691.
- Fischbacher, U. (2007). z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2), 171-178.
- Galizzi, M. M., & Navarro-Martínez, D. (2019). On the external validity of social preference games: a systematic lab-field study. *Management Science*, 65(3), 976-1002.

Golman, R., Bhatia, S. & Kane, P. (in press). The dual accumulator model of strategic deliberation and decision making. *Psychological Review*.

Greiner, B. (2015). Subject pool recruitment procedures: organizing experiments with ORSEE. *Journal of the Economic Science Association*, 1(1), 114-125.

Harrison, G. W. (1989). Theory and misbehavior of first-price auctions. *American Economic Review*, 749-762.

Hristova, E., & Grinberg, M. (2005). Information acquisition in the iterated prisoner's dilemma game: An eye-tracking study. In *Proceedings of the 27th annual conference of the cognitive science society* (pp. 983-988). Hillsdale, NJ: Lawrence Erlbaum.

Ho, T. H., Camerer, C., & Weigelt, K. (1998). Iterated dominance and iterated best response in experimental "p-beauty contests". *American Economic Review*, 88(4), 947-969.

Hoffmann, T. (2014). The Effect of Belief Elicitation Game Play. Working paper.

Iriberri, N., & Rey-Biel, P. (2013). Elicited beliefs and social information in modified dictator games: What do dictators believe other dictators do?. *Quantitative Economics*, 4(3), 515-547.

Ivanov, A. (2011). Attitudes to ambiguity in one-shot normal-form games: An experimental study. *Games and Economic Behavior*, 71(2), 366-394.

McKelvey, R. D., & Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1), 6-38.

Nagel, R. (1995). Unraveling in guessing games: An experimental study. *American Economic Review*, 85(5), 1313-1326.

Nyarko, Y., & Schotter, A. (2002). An experimental study of belief learning using elicited beliefs. *Econometrica*, 70(3), 971-1005.

Polonio, L., Di Guida, S., & Coricelli, G. (2015). Strategic sophistication and attention in games: An eye-tracking study. *Games and Economic Behavior*, 94, 80-96.

Polonio, L., & Coricelli, G. (2019). Testing the level of consistency between choices and beliefs in games using eye-tracking. *Games and Economic Behavior*, 113, 566-586.

Pulford, B. D., Colman, A. M., & Loomes, G. (2018). Incentive magnitude effects in experimental games: Bigger is not necessarily better. *Games*, 9(1), 4.

Rey-Biel, P. (2009). Equilibrium play and best response to (stated) beliefs in normal form games. *Games and Economic Behavior*, 65(2), 572-585.

Rydval, O., Ortmann, A., & Ostadnick, M. (2009). Three very simple games and what it takes to solve them. *Journal of Economic Behavior & Organization*, 72(1), 589-601.

Schlag, K. H., Tremewan, J., & Van der Weele, J. J. (2015). A penny for your thoughts: A survey of methods for eliciting beliefs. *Experimental Economics*, 18(3), 457-490.

Schotter, A., & Trevino, I. (2014). Belief elicitation in the laboratory. *Annual Review of Economics*, 6, 103-128.

Sobel, J. (2005). Interdependent preferences and reciprocity. *Journal of Economic literature*, 43(2), 392-436.

Stewart, N., Gächter, S., Noguchi, T., & Mullett, T. L. (2016). Eye movements in strategic choice. *Journal of behavioral decision making*, 29(2-3), 137-156.

Stahl, D. O., & Wilson, P. W. (1994). Experimental evidence on players' models of other players. *Journal of Economic Behavior & Organization*, 25(3), 309-327.

Stahl, D. O., & Wilson, P. W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1), 218-254.

Sutter, M., Czermak, S., & Feri, F. (2013). Strategic sophistication of individuals and teams. Experimental evidence. *European Economic Review*, 64, 395-410.

Zonca, J., Coricelli, G., & Polonio, L. (2018). Disclosing the link between cognitive reflection and sophistication in strategic interaction: the crucial role of game representation. Working paper.

Appendix A

Additional Tables and Figures

Table A1. Consistency of beliefs and choices of the row (BLUE) player with theoretical predictions

Game	Choice task			Belief task		
	Nash	Level-1	Level-2	Nash	Level-1	Level-2
1	0.38	0.46	0.38	0.26	0.55	0.26
2	0.30	0.58	0.58	0.27	0.54	0.19
3	0.38	0.50	0.38	0.27	0.54	0.27
4	0.30	0.48	0.48	0.27	0.31	0.42
5	0.12	0.74	0.74	0.17	0.66	0.17
6	0.56	0.40	0.04	0.52	0.34	0.34
7	0.06	0.76	0.76	0.16	0.60	0.23
8	0.10	0.76	0.14	0.33	0.45	0.22
Average	0.27	0.58	0.44	0.29	0.50	0.26

Notes: The left side shows the average proportion of choices in accordance with the different models' predictions. The right side shows the stated beliefs (average probabilities on model predictions) about the model of choice of the counterpart. For Game 2, which has two Nash Equilibria, we show the rate with which subjects chose the Pareto dominant Nash Equilibrium. The total average for the Nash equilibrium is calculated over the games with unique equilibrium (excluding Game 2).

Table A2. Consistency of beliefs and choices of the column (RED) player with theoretical predictions

Game	Choice task			Belief task		
	Nash	Level-1	Level-2	Nash	Level-1	Level-2
1	0.24	0.60	0.24	0.38	0.38	0.38
2	0.36	0.30	0.34	0.24	0.58	0.58
3	0.46	0.34	0.46	0.24	0.51	0.24
4	0.22	0.08	0.70	0.31	0.54	0.54
5	0.22	0.52	0.26	0.12	0.63	0.63
6	0.42	0.50	0.50	0.49	0.46	0.05
7	0.12	0.48	0.40	0.08	0.65	0.65
8	0.20	0.58	0.22	0.25	0.52	0.23
Average	0.27	0.43	0.39	0.27	0.53	0.41

Notes: The left side shows the average proportion of choices in accordance with the different models' predictions. The right side shows the stated beliefs (average probabilities on model predictions) about the model of choice of the counterpart. For Game 2, which has two Nash Equilibria, we show the rate with which subjects chose the Pareto dominant Nash Equilibrium. The total average for the Nash equilibrium is calculated over the games with unique equilibrium (excluding Game 2).

Table A3. Mean normalized ranking score for pairs of own and other's payoffs

Own payoff	Other's payoff					
	0	2	4	6	8	10
0	0.96	0.89	0.90	0.85	0.87	0.81
2	0.77	0.61	0.75	0.72	0.74	-
4	0.55	0.54	0.36	0.58	0.65	0.52
6	0.26	0.29	0.34	0.15	0.41	-
8	0.28	0.18	0.31	0.18	0.09	-
10	0.12	-	0.10	-	-	0.02

Notes: Scores are calculated as follows: $x' = \frac{x - \min(x)}{\max(x) - \min(x)}$, where x is the actual rank given, and $\min(x)$ and $\max(x)$ are the minimum and maximum rank possible in a game. In particular, the maximum rank possible differs between 7 and 9, depending on the number of unique payoff pairs. The normalized score ranges from 0 to 1, where lower scores correspond to more preferred (lower-ranked) payoff pairs.


Table A4. Consistency of choices per game and player type using expected payoffs and expected rankings.

	All players		Selfish		Inequity Averse	
	Exp. payoffs	Exp. ranking	Exp. payoffs	Exp. ranking	Exp. payoffs	Exp. ranking
1	0.44	0.52	0.47	0.44	0.41	0.61
2	0.54	0.52	0.49	0.45	0.59	0.59
3	0.49	0.49	0.53	0.51	0.45	0.48
4	0.55	0.54	0.49	0.51	0.59	0.57
5	0.56	0.65	0.58	0.67	0.52	0.61
6	0.68	0.68	0.65	0.65	0.70	0.70
7	0.68	0.70	0.69	0.69	0.66	0.70
8	0.42	0.53	0.47	0.49	0.36	0.57
Average	0.54	0.58	0.55	0.55	0.54	0.61

Table A5. Regression analysis of response times

Dependent variable:	Response time	
	(1)	(2)
Std. dev. in expected payoffs	-2.978*** (0.680)	
Std. dev. in expected rankings		-4.357*** (0.094)
Constant	26.854*** (2.239)	27.875*** (2.151)
# Observations (clusters)	800 (100)	800 (100)
(Pseudo) R ²	0.022	0.022

Notes: This table reports coefficient estimates from OLS regressions. The dependent variable is how much time subjects spend on choosing their strategy. As independent variables we use the standard deviation in expected payoffs or expected rankings across the three available strategies for each player. Standard errors clustered at the individual level are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Figure A1. Distribution of the number of optimal choices per subject using expected payoffs (left panel) and expected rankings (right panel)

Appendix B

Comparing Structured vs. Unstructured Responses

Table B1. Comparison of game play based on different models' prediction.

Game	Nash			Level-1			Level-2		
	S	U	<i>p-value</i>	S	U	<i>p-value</i>	S	U	<i>p-value</i>
1	0.31	0.33	0.768	0.53	0.51	0.788	0.31	0.33	0.768
2	0.33	0.43	0.172	0.44	0.31	0.061	0.46	0.35	0.125
3	0.42	0.36	0.407	0.42	0.41	0.943	0.42	0.36	0.407
4	0.26	0.34	0.224	0.28	0.32	0.553	0.59	0.54	0.506
5	0.17	0.17	0.997	0.63	0.54	0.218	0.50	0.49	0.883
6	0.49	0.59	0.186	0.45	0.39	0.428	0.27	0.23	0.566
7	0.09	0.16	0.147	0.62	0.56	0.428	0.58	0.60	0.824
8	0.15	0.23	0.140	0.67	0.61	0.358	0.18	0.16	0.706
Average	0.27	0.31	0.139	0.50	0.46	0.162	0.41	0.38	0.348

Notes: Average proportion of choices in accordance with the different models' predictions in the Structured (S) and the Unstructured (U) treatment. P-values from logistic regressions with robust standard errors (clustered at the individual level). For Game 2, which has two Nash Equilibria, we show the rate with which subjects chose the Pareto dominant Nash Equilibrium. The total average for the Nash equilibrium is calculated over the games with unique equilibrium (excluding Game 2).

Table B2. Comparison of game play of the row (BLUE) player based on different models' prediction.

Game	Nash			Level-1			Level-2		
	S	U	<i>p-value</i>	S	U	<i>p-value</i>	S	U	<i>p-value</i>
1	0.38	0.28	0.283	0.46	0.51	0.620	0.38	0.28	0.283
2	0.30	0.32	0.839	0.58	0.43	0.132	0.58	0.43	0.132
3	0.38	0.30	0.397	0.50	0.47	0.755	0.38	0.30	0.397
4	0.30	0.34	0.671	0.48	0.53	0.611	0.48	0.53	0.611
5	0.12	0.17	0.486	0.74	0.62	0.199	0.74	0.62	0.199
6	0.56	0.60	0.723	0.40	0.36	0.700	0.04	0.04	0.950
7	0.06	0.11	0.415	0.76	0.77	0.945	0.76	0.77	0.945
8	0.10	0.15	0.469	0.76	0.70	0.523	0.14	0.15	0.901
Average	0.27	0.28	0.883	0.58	0.55	0.416	0.44	0.39	0.250

Notes: Average proportion of choices in accordance with the different models' predictions in the Structured (S) and the Unstructured (U) treatment. P-values from logistic regressions with robust standard errors (clustered at the individual level). For Game 2, which has two Nash Equilibria, we show the rate with which subjects chose the Pareto dominant Nash Equilibrium. The total average for the Nash equilibrium is calculated over the games with unique equilibrium (excluding Game 2).

Table B3. Comparison of game play of the column (RED) player based on different models' prediction.

Game	Nash			Level-1			Level-2		
	S	U	<i>p-value</i>	S	U	<i>p-value</i>	S	U	<i>p-value</i>
1	0.24	0.38	0.132	0.60	0.51	0.379	0.24	0.38	0.132
2	0.36	0.53	0.092	0.30	0.19	0.221	0.34	0.28	0.502
3	0.46	0.43	0.734	0.34	0.36	0.824	0.46	0.43	0.734
4	0.22	0.34	0.191	0.08	0.11	0.657	0.70	0.55	0.139
5	0.22	0.17	0.540	0.52	0.47	0.611	0.26	0.36	0.283
6	0.42	0.57	0.132	0.50	0.43	0.465	0.50	0.43	0.465
7	0.12	0.21	0.226	0.48	0.36	0.242	0.40	0.43	0.800
8	0.20	0.32	0.185	0.58	0.51	0.495	0.22	0.17	0.540
Average	0.27	0.35	0.074	0.43	0.37	0.189	0.39	0.38	0.804

Notes: Average proportion of choices in accordance with the different models' predictions in the Structured (S) and the Unstructured (U) treatment. P-values from logistic regressions with robust standard errors (clustered at the individual level). For Game 2, which has two Nash Equilibria, we show the rate with which subjects chose the Pareto dominant Nash Equilibrium. The total average for the Nash equilibrium is calculated over the games with unique equilibrium (excluding Game 2).

Appendix C

Experimental Instructions

STRUCTURED TREATMENT

INTRODUCTION

Welcome. You are now participating in a study about decision making. If you follow the instructions carefully you might earn a considerable amount of money which will be paid at the end of the experiment in private and in cash. It is important that during the experiment you remain silent. If you have questions or need assistance, please raise your hand. A member of the experimental team will come to you and answer them in private.

The experiment consists of **one task**, which will be described to you below in detail. You will be paired with someone else in the room. Half of the people in the room will be given the role of **BLUE** decision makers and the other half will be given the role of **RED** decision maker.

You have randomly selected to be a (BLUE or RED) decision maker. You will be paired at random with someone of the other colour. Please remember this as you will stay in that role during the whole experiment.

At the end of the experiment, after everyone has made all of their decisions, the computer will randomly pick **one decision** for each pair from the task and we will pay you according to the decisions that were made by you and the person you are paired with. The amount you will get from the randomly selected decision will be added to the £3 'show-up' fee.

Important: You must think carefully about each of your decisions as each of them can determine your earnings from this experiment. All decisions are anonymous, i.e., you will never be told the identity of the person you are paired with - it is equally likely to be anyone of the other colour.

THE TASK:

What follows is an *example*, intended to give you a chance to *practice* and *check your understanding* before doing any tasks for real.

Please click on OK to proceed.

THE TASK

GENERAL DECISION SITUATION

In the course of the task you will see a series of different situations in which you will interact with another person of the other colour. In each situation you will be asked to make various decisions. To familiarise you with the idea, please look at the decision scenario below which serves as an **example**.

In the grid both players, **BLUE** and **RED**, have to decide between three options. The **BLUE** decision maker can choose between the **rows** labelled **A**, **B** and **C**, while the **RED** decision maker chooses between the **columns** labelled **D**, **E** and **F**. The payments that the **BLUE** decision maker might receive are coloured **BLUE** and the possible payments for the **RED** decision-maker are shown coloured **RED**.

	D	E	F
A	4, 4	8, 6	0, 4
B	6, 8	0, 0	4, 10
C	4, 0	10, 4	2, 2

How much each person will get paid depends on the choices each makes and how these choices interact. For example, suppose the **BLUE** decision maker chooses **C** and the **RED** decision maker chooses **D**. By looking at the pair of numbers where **C** and **D** intersect, we see that **BLUE** gets a payment of 4 and **RED** gets a payment of 0. Now suppose that **BLUE** chooses **C**, as before, but **RED** chooses **F** instead of **D**. In this case, **BLUE** gets 2 and **RED** gets 2. The nine boxes in the grid show the nine pairs of payment that are possible, depending on the choices each person makes.

As a check that everyone has understood, please type in answers to the following questions:

If **BLUE** chooses **A** and **RED** chooses **F**, **BLUE** will get ____ and **RED** will get ____

If **BLUE** chooses **B** and **RED** chooses **D**, **BLUE** will get ____ and **RED** will get ____

When you have typed in your answers, please click on OK.

MAKING DECISIONS

In the task we ask you to make **three types of decisions** which we will explain to you with the help of the example from before.

	D	E	F
A	4, 4	8, 6	0, 4
B	6, 8	0, 0	4, 10
C	4, 0	10, 4	2, 2

In the **first type of decision**, you have to rank the number of possible pairs of payoffs from your most preferred pair to your least preferred pair. Each pair consists of a money amount for the **BLUE** person

and an amount for the **RED** person. For example, one pair states that **BLUE** receives an amount of **4** and **RED** receives an amount of **4**. Another option states that **BLUE** receives an amount of **6** and **RED** receives an amount of **8**. In total the number of pairs of payoffs you have to rank is equal to the number of **unique** combinations you see in a specific situation. You are asked to rank-order these options according to the pair of payoffs you personally prefer most, second, third, and so on down to the pair you like least.

The following table shows an example.

Option	Rank from 1 to 9
4, 4	
8, 6	
0, 4	
6, 8	
0, 0	
4, 10	
4, 0	
10, 4	
2, 2	

In this example, someone who feels that the best result would be **BLUE** getting £4 and **RED** getting £4 would type in a 1 next to **4, 4** to show that they rate this pair as their 1st-ranked pair of payoffs. **(You may not agree with that yourself, but this is only an example.)** Then if their 2nd-ranked outcome is for **BLUE** getting **6** and **RED** getting **8**, they would type the number 2 next to that pair **6, 8**. And so on, until all nine pairs have been ranked from most preferred (1) down to least preferred (9).

Remember that the first number in each option is relevant for the **BLUE person, whereas the second number is relevant for the **RED** person.**

So now, in order to give you some practice, look at the grid and type in your ranks from your most preferred (1) down to your least preferred (9). When you have typed in your answers, please click on OK.

In the **second type of decision** we ask you to estimate what the participants of the other colour will actually choose. As you know, there are 10 people in this session who will be **BLUE** decision makers and another 10 people who will be **RED** decision makers. We will ask you to give your best judgment about which of the three alternatives the 10 **BLUE** (**RED**) decision makers will choose, i.e., how many (if any) of them will choose **A** (**D**), how many (if any) will choose **B** (**E**), and how many (if any) will choose **C** (**F**). So now, in order to give you some practice, look at the grid and type in your estimates below:

My best estimate of how many of the BLUE (RED) decision makers will choose A (D) is _____

My best estimate of how many of the BLUE (RED) decision makers will choose B (E) is _____

My best estimate of how many of the BLUE (RED) decision makers will choose C (F) is _____

Please make sure that your three estimates add up to 10.

In the **third type of decision** we finally ask you to simply choose YOUR option. In the light of what you think BLUE (RED) might do and what you want to achieve, which option – D, E or F (A, B, or C) – do you choose?

So now, in order to give you some practice, look at the grid and make a decision by selecting one of the three letters below. When you have typed in your answers, please click on OK.

Calculation of earnings

Your earnings for the task will be determined as follows. First, the computer will **randomly select one out of the series of decision situations**. Then, the computer will randomly decide whether you are paid according to your **decision of type one**, your **decision of type two**, or your **decision of type three**. Below we explain each of these scenarios in more detail.

If your **decision of type one** is selected to determine your payment, this is what would happen: At the end of the experiment, we will randomly pick two of the pairs of payoffs from the selected decision scenario. We will then randomly determine whether each pair of people gets paid according to how the BLUE person ranked those two pairs or according to how the RED person ranked them. We will then check which option out of the two the selected person ranked higher, and pay you and the other player the amount stated in that option. For example if your decision is selected to be implemented and you ranked 6, 8 higher than 4, 4, then the BLUE decision maker will receive an amount of 6 and the RED decision maker will receive an amount of 8. Alternatively, if your decision is selected to be implemented and you ranked 4, 4 higher than 6, 8, then the BLUE decision maker will receive an amount of 4 and the RED decision maker will receive an amount of 4.

If your **decision of type two** is selected to determine your payment, this is what would happen: First, we would pick one of the three letters A, B, or C (D, E, or F) at random. Then, we would compare your estimate *for that particular letter* with what the 10 BLUE (RED) decision makers actually did in the randomly selected decision scenario. If your estimate for that letter is correct, you get £5. If your estimate is wrong, you get £0. There is no prize for being ‘close’ – you either get it right or you don’t, so please think carefully.

If your **decision of type three** is selected to determine your payment, this is what would happen: First, we would look up your type three decision in the randomly selected scenario. Second, we would look up the corresponding type three decision of the BLUE/RED participant you are matched with in the randomly selected decision scenario. We would then look up the pair of numbers where the two decisions intersect and pay player BLUE and RED the BLUE and RED amount, respectively.

Please click on OK to start the task. Notice that for the whole duration of task, whenever you press OK you will not be able to return to your previous choice, so please think carefully about each of your decisions before pressing the OK button.

UNSTRUCTURED TREATMENT

INTRODUCTION

Welcome. You are now participating in a study about decision making. If you follow the instructions carefully you might earn a considerable amount of money which will be paid at the end of the experiment in private and in cash. It is important that during the experiment you remain silent. If you have questions or need assistance, please raise your hand. A member of the experimental team will come to you and answer them in private.

The experiment consists of **one task**, which will be described to you below in detail. You will be paired with someone else in the room. Half of the people in the room will be given the role of **BLUE** decision makers and the other half will be given the role of **RED** decision maker.

You have randomly selected to be a (BLUE or RED) decision maker. You will be paired at random with someone of the other colour. Please remember this as you will stay in that role during the whole experiment.

At the end of the experiment, after everyone has made all of their decisions, the computer will randomly pick **one decision** for each pair from the task and we will pay you according to the decisions that were made by you and the person you are paired with. The amount you will get from the randomly selected decision will be added to the £3 'show-up' fee.

Important: You must think carefully about each of your decisions as each of them can determine your earnings from this experiment. All decisions are anonymous, i.e., you will never be told the identity of the person you are paired with - it is equally likely to be anyone of the other colour.

THE TASK:

What follows is an *example*, intended to give you a chance to *practice* and *check your understanding* before doing any tasks for real.

Please click on OK to proceed.

THE TASK

GENERAL DECISION SITUATION

In the course of the task you will see a series of different situations in which you will interact with another person of the other colour. In each situation you will be asked to make various decisions. To familiarise you with the idea, please look at the decision scenario below which serves as an **example**.

In the grid both players, **BLUE** and **RED**, have to decide between three options. The **BLUE** decision maker can choose between the **rows** labelled **A**, **B** and **C**, while the **RED** decision maker chooses between the **columns** labelled **D**, **E** and **F**. The payments that the **BLUE** decision maker might receive are coloured **BLUE** and the possible payments for the **RED** decision-maker are shown coloured **RED**.

	D	E	F
A	4, 4	8, 6	0, 4
B	6, 8	0, 0	4, 10
C	4, 0	10, 4	2, 2

How much each person will get paid depends on the choices each makes and how these choices interact. For example, suppose the **BLUE** decision maker chooses **C** and the **RED** decision maker chooses **D**. By looking at the pair of numbers where **C** and **D** intersect, we see that **BLUE** gets a payment of 4 and **RED** gets a payment of 0. Now suppose that **BLUE** chooses **C**, as before, but **RED** chooses **F** instead of **D**. In this case, **BLUE** gets 2 and **RED** gets 2. The nine boxes in the grid show the nine pairs of payment that are possible, depending on the choices each person makes.

As a check that everyone has understood, please type in answers to the following questions:

If **BLUE** chooses **A** and **RED** chooses **F**, **BLUE** will get ____ and **RED** will get ____

If **BLUE** chooses **B** and **RED** chooses **D**, **BLUE** will get ____ and **RED** will get ____

When you have typed in your answers, please click on OK.

MAKING DECISIONS

In the task we ask you to make **decisions** which we will explain to you with the help of the example from before.

	D	E	F
A	4, 4	8, 6	0, 4
B	6, 8	0, 0	4, 10
C	4, 0	10, 4	2, 2

For each possible decision situation we ask you to simply choose YOUR option. In the light of what you think **BLUE** (**RED**) might do and what you want to achieve, which option – **D**, **E** or **F** (**A**, **B**, or **C**) – do you choose? Type the letter of your choice in the box below and then click OK.

So now, in order to give you some practice, look at the grid and make a decision by selecting one of the three letters below. When you have typed in your answers, please click on OK.

Calculation of earnings

Your earnings for the task will be determined as follows. First, the computer will **randomly select one out of the series of decision situations**.

First, we will look up your decision at the randomly selected situation. Second, we would look up the corresponding decision of the participant of the other colour you are matched with in the randomly selected decision situation. We would then look up the pair of numbers where the two decisions intersect and pay player **BLUE** and **RED** the **BLUE** and **RED** amount, respectively.

Please click on OK to start the task. Notice that for the whole duration of task, whenever you press OK you will not be able to return to your previous choice, so please think carefully about each of your decisions before pressing the OK button.