



Investigating the failure to best respond in experimental games

Despoina Alempaki¹ · Andrew M. Colman² · Felix Kölle^{3,4} · Graham Loomes¹ · Briony D. Pulford²

Received: 14 December 2020 / Revised: 15 June 2021 / Accepted: 15 June 2021 /
Published online: 1 July 2021
© The Author(s) 2021

Abstract

We examine strategic sophistication using eight two-person 3×3 one-shot games. To facilitate strategic thinking, we design a ‘structured’ environment where subjects first assign subjective values to the payoff pairs and state their beliefs about their counterparts’ probable strategies, before selecting their own strategies in light of those deliberations. Our results show that a majority of strategy choices are inconsistent with the equilibrium prediction, and that only just over half of strategy choices constitute best responses to subjects’ stated beliefs. Allowing for other-regarding considerations increases best responding significantly, but the increase is rather small. We further compare patterns of strategies with those made in an ‘unstructured’ environment in which subjects are not specifically directed to think strategically. Our data suggest that structuring the pre-decision deliberation process does not affect strategic sophistication.

Keywords Equilibrium · Best response · Strategic thinking · Social preferences · Beliefs

JEL Classification A13 · C72 · C91 · C92 · D84

✉ Despoina Alempaki
despoina.alempaki@wbs.ac.uk

✉ Felix Kölle
felix.koelle@uni-koeln.de

¹ Warwick Business School, University of Warwick, Coventry CV4 7AL, UK

² Department of Neuroscience, Psychology and Behaviour, University of Leicester, Leicester LE1 7RH, UK

³ Centre for Decision Research and Experimental Economics, School of Economics, University of Nottingham, Nottingham, UK

⁴ Faculty of Management, Economics and Social Sciences, University of Cologne, Cologne, Germany

1 Introduction

The experimental game literature has produced a number of studies showing that a substantial proportion of individuals' strategy choices neither correspond to equilibrium predictions, nor are 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 Camerer, 2003; Costa-Gomes & Weizsäcker, 2008; Danz et al., 2012; Hoffmann, 2014; Polonio & Coricelli, 2019; Sutter et al., 2013).

Various possibilities have been suggested to explain deviations from equilibrium behaviour, including 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 strategies, outcomes and payoffs (e.g., Bosch-Rosa & Meissner, 2020; Cason & Plott, 2014; Chou et al., 2009; Cox & James, 2012; Fehr & Huck, 2016; Rydval et al., 2009; Zonca et al., 2020). Studies using choice process data (e.g., Brocas et al., 2014; Devetag et al., 2016; Hristova & Grinberg, 2005; Polonio et al., 2015; Stewart et al., 2016) suggest that when choosing a strategy, 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 payoffs, 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 reason strategically might lie in the existence of other-regarding motives (see e.g., Cooper & Kagel, 2016; Fehr & Schmidt, 2006; Sobel, 2005 for overviews of the literature). If individuals' strategy choices are not solely driven by self-interest but involve social preferences, it should not be surprising that subjects depart from optimal behaviour 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 & Coricelli, 2019; Polonio et al., 2015). While evidence from this literature suggests that other-regarding motives may interact with the observed levels of strategic sophistication, the correspondence between strategy choices 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.

In this paper, we focus primarily on two questions. First, can we increase best response rates by encouraging subjects to think more systematically about the

subjective values they assign to the payoff pairs and about the weights they attach to other players' possible actions? Second, will allowance for other-regarding considerations increase best response rates?

We try to answer these questions by conducting a laboratory experiment in which subjects are presented with a set of eight one-shot two-person 3×3 normal-form games.¹ To encourage strategic sophistication, we design a *Structured* treatment, where we prompt players to think about their evaluation of payoff pairs in conjunction with their beliefs about other players' possible strategy 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 think about the subjective value they attach to the various cells in each 3×3 game. The joint evaluation of payoffs requires subjects to attend more closely to their counterparts' incentives and motives in general. This increased attention towards others might translate into a higher degree of strategic thinking by helping less sophisticated subjects to become better at predicting their counterparts' strategies.² At the same time, based on subjects' stated rankings, we may infer something about their other-regarding considerations, thus relaxing the assumption of pure self-interest (as used by much of the literature) in cases where it does not seem appropriate. Some previous studies (e.g., Bayer & 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-Martínez, 2019, for evidence and a meta-analysis). By contrast, our ranking task explores the expression of at least some other-regarding preferences in a format *intrinsic* to the actual games under consideration. Thus, the purpose of the joint evaluation of payoffs is two-fold: to increase strategic sophistication, while providing some measure of other-regarding preferences.³

Second, using a *belief-elicitation* 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

¹ All data and code supporting the findings of the experiment are available from the Open Science Framework, accessible at <https://osf.io/zhufe/>.

² The joint evaluation feature of the ranking task is of particular importance given previous evidence suggesting that a non-negligible fraction of subjects neglect their counterpart's payoffs in strategic environments or focus only on a subset of all possible game outcomes (e.g., Devetag et al., 2016; Polonio et al., 2015). See also the discussion in Weizsäcker (2003), where adding context to the usually abstract matrix representations of normal-form games is suggested as a way to increase attention to others' incentives.

³ As discussed later, we do not claim that this task will necessarily capture *all* possible components of social preferences. However, to the extent that players take account of other players' payoffs and the relative values between their own and others' payoffs, the ranking task prompts them to weigh up the trade-offs involved.

has obtained mixed results: some studies find that belief elicitation does not affect game play (e.g., Costa-Gomes & Weizsäcker, 2008; Ivanov, 2011; Nyarko & Schotter, 2002; Polonio & Coricelli, 2019), whereas others show that whether belief elicitation influences game play depends on the properties of the game (e.g., Hoffmann, 2014; see also Schlag et al., 2015 and Schotter & 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. As such, the addition of the ranking task might influence consistency rates as well as game play.

Our main results can be summarized as follows. In line with previous literature (e.g., Costa-Gomes & Weizsäcker, 2008; Danz et al., 2012; Hoffmann, 2014; Polonio & Coricelli, 2019; Sutter et al., 2013) we find that a sizeable proportion of players do not choose equilibrium strategies, and fail to best respond to their own stated beliefs. Importantly, while we find that the rate of best response increases significantly when we allow for other-regarding preferences as revealed by the ranking task, the overall magnitude of the difference is relatively small (54% vs. 57%)—although for the subset of subjects we classify as inequity averse, best response rates increase from 54% to 61% when using rankings instead of own payoffs. This finding reveals that the observed increase at the aggregate level is almost entirely driven by those subjects who are not motivated only by their own earnings. Still, even for this subset, a significant share of non-optimal play remains, which is line with evidence from some previous studies showing that, in contexts different from ours, other-regarding preferences alone cannot fully explain deviations from equilibrium behaviour (e.g., Andreoni, 1995; Andreoni & Blanchard, 2006; Fischbacher et al., 2009; Krawczyk & Le Lec, 2015).

Our results also suggest that the likelihood of choosing non-optimally is decreasing in the costs of doing so. That is, while non-optimal strategy 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 & Palfrey, 1995) or some other ‘error’ model (see, for instance, Harrison, 1989, for a discussion of the ‘flat maximum’ critique).

To examine more systematically how far and in what ways our intervention alters strategic sophistication, if at all, we compare the patterns of strategy choice in 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 eight one-shot 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 the patterns of choice might be affected. We find overall no impact of structuring the decision process: our attempts to induce ‘harder thinking’ does not affect chosen strategies.

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

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

Fig. 1 Experimental games. Underlined payoffs indicate the Nash equilibria in pure strategies. All payoffs are in UK pounds

2 The experiment

We chose a set of eight one-shot 3 × 3 normal-form games that were adjusted versions of the games used in Colman et al. (2014). The games are displayed in Fig. 1.⁴ 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. All payoffs are 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 chosen strategies in these games suggests relatively low-effort thinking, which gives us enough room to detect whether structuring responses leads to higher levels of strategic 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 & Crawford, 2006; Ho et al., 1998; Nagel, 1995; Stahl & Wilson, 1994, 1995) to allow for bounded depth of reasoning. The *Level-1* model assumes that a player best responds

⁴ Our games differ from Colman et al. (2014) as follows. We doubled all the original payoffs in order to bring earnings more in 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.

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

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

to the belief that assigns uniform probabilities to their counterpart's 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. Subjects 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' responses until the end of the experiment.

For each game in the Structured treatment, subjects completed three different tasks: a *ranking*, a *belief-elicitation* and a *strategy-choice* task.⁵ 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, one out of the eight games was randomly selected for payment, and another random draw directed whether subjects' earnings were determined according to the ranking, the belief-elicitation, or the strategy-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. 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

⁵ In the experiment, the different tasks were called Type I, Type II, and Type III decisions respectively (see Figures 2–4, as well as the experimental instructions in Online Appendix C).

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

Fig. 2 Screenshot of the Structured treatment's ranking task

appear in the game from top left to bottom right.⁶ They stated their ranking of these pairs by typing in a number between 1 and 9, where 1 corresponded to the pair they liked best and 9 indicated their least preferred pair.

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. For example, if the randomly selected player was the BLUE player, and the two randomly selected payoff pairs were (8, 6) and (10, 4), we looked at which payoff pair the BLUE player ranked higher and paid both players accordingly (either £8 for BLUE and £6 for RED, or £10 for BLUE and £4 for RED player).

Once subjects had submitted their rankings, they proceeded automatically to the *belief-elicitation* 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.⁷ A screenshot of the interface used in the belief-elicitation task is shown in Fig. 3. In the event that the belief-elicitation 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 strategy choice with

⁶ 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 *strategy-choice* task.

⁷ All sessions in the Structured treatment were conducted with 20 participants each, 10 BLUE and 10 RED players.

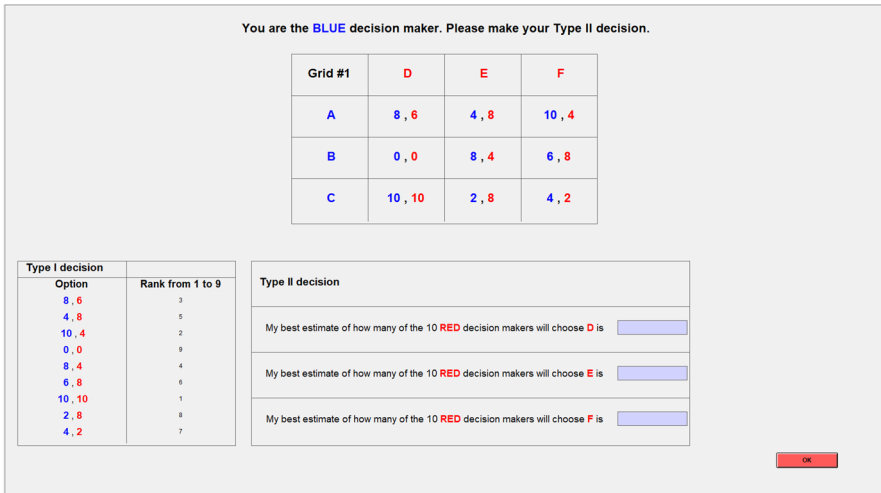


Fig. 3 Screenshot of the Structured treatment's belief-elicitation task

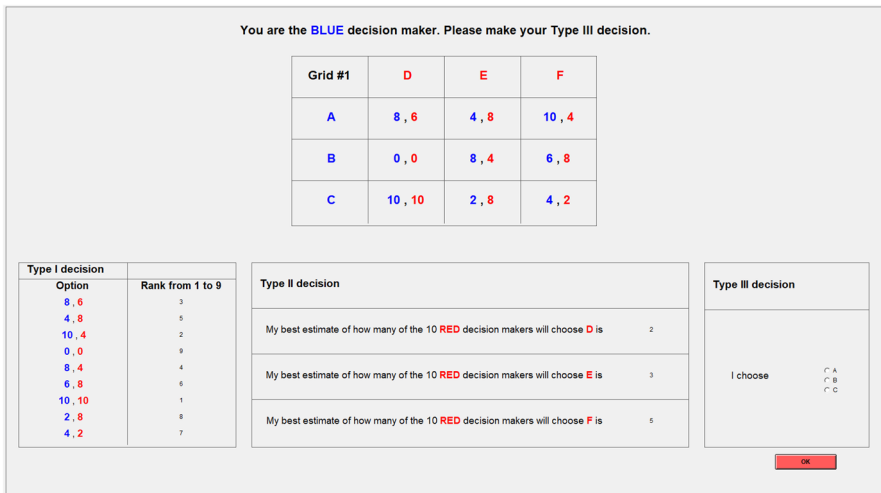


Fig. 4 Screenshot of the Structured treatment's strategy-choice task

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.⁸

⁸ 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).

Finally, on the last screen of each game subjects had to indicate their preferred strategy in the *strategy-choice* task (see Fig. 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 chosen strategies. For instance, if the BLUE player chose strategy A, and the RED player strategy D, the BLUE player would receive £8 and the RED player £6.

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-elicitation tasks (as shown in Fig. 4), making it as easy as we could for them to think strategically in light of their previous deliberations, if that is what they wished to do.

In order to measure whether structuring subjects' decision-making process and inducing them to think about all relevant aspects of a game before making their decisions had any substantial systematic effect on strategy choice at all, we ran a separate Unstructured control treatment. Using the exact same eight one-shot games, subjects were simply asked to state—without any prior deliberation tasks—which of the three available strategies they wanted to play. In this unstructured environment, 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 not have affected their strategy choices when they saw each game for the first time. To avoid any income effects on subjects' behaviour in the repetitions of the Unstructured treatment, we randomly selected for payment a game from any sequence, and not necessarily from the first. We then paid subjects according to the intersection of their chosen strategies.⁹

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, subjects were asked to read some preliminary instructions of an example 3×3 game and to demonstrate they fully understand the required tasks by answering a series of control questions before they could proceed to the actual experiment.

⁹ The number of repetitions in the Unstructured treatment was determined such that both the Structured and Unstructured treatments are of similar duration, so expected hourly earnings do not differ.

Table 2 Average proportion of chosen strategies and stated beliefs on models' predictions

Game	Strategy-choice task			Belief-elicitation 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

The left side shows the average proportion of strategy choices in accordance with the different models' predictions. The right side shows the average probability with which subjects estimated each of the different models' predictions to be chosen by their counterpart. For Game 2, which has two Nash equilibria, we show the rate with which subjects chose the Pareto dominant Nash equilibrium, and the rate with which they estimate their counterpart will choose the Pareto dominant Nash equilibrium. Tables A1 and A2 in Online 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)

3 Results

We organise our results as follows. In Sect. 3.1, we start with a descriptive analysis of chosen strategies and stated beliefs in the Structured treatment, and discuss whether behaviour in this environment exhibits high degree of strategic sophistication, assuming subjects are only interested in maximizing their own payoff. In Sect. 3.2, we turn our focus 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 strategies. In Sect. 3.3, we discuss possible determinants of non-optimal behaviour. Finally, in Sect. 3.4 we compare patterns of chosen strategies across the Structured and the Unstructured treatments.

3.1 Strategy choices, beliefs and best responses

We start our analysis by focusing on the average proportion of chosen strategies that were predicted by the Nash equilibrium together with the Level-1 and Level-2 models. This is shown in the left side of Table 2 (see the Strategy-choice task columns). On average, the proportion of Nash equilibrium strategy 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 strategies in games with a unique Nash equilibrium

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

(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 & Weizsäcker, 2008; Polonio & Coricelli, 2019), the Level-1 and Level-2 models both outperform equilibrium predictions. The fractions of chosen strategies 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.

Next, we contrast the patterns of chosen strategies with the average probability with which subjects estimated that the counterpart is playing each model's predictions. As depicted by the right side of Table 2 (see the Belief-elicitation task columns in Table 2), for stated beliefs we observe a similar pattern as for chosen strategies: Most subjects do not expect their counterparts to play according to equilibrium: the average probability with which subjects estimated that their counterparts are playing the equilibrium strategy (in games with a unique Nash equilibrium) amounts to only 0.28. Instead, following the same pattern as the chosen strategies, subjects most often predicted their counterparts would play according to the Level-1 model (probability=0.52), followed by the Level-2 model (probability=0.34).

The fact that most subjects expect that their counterpart will play their Level-1 strategy, and at the same time, as a response, they also choose their Level-1 strategy, already provides a first indication at the aggregate level that subjects often did not best respond. If they were best responding to their own stated beliefs, they should have moved one step up in the hierarchy and chosen strategies predicted by the Level-2 model. This is not what we observe in our data, since the proportion of Level-2 chosen strategies (probability=0.41) is lower than the proportion of Level-1 beliefs (probability=0.52) (see e.g., Polonio & Coricelli, 2019 for similar evidence).

To provide more in-depth evidence on subjects' best response behaviour, in the following, we investigate the level of consistency between chosen strategies 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 those 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

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

Ranking from 1 (most preferred) to 9 (least preferred). 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 Online Appendix A we display an alternative version of Table 3 in which we normalize the rankings. The results remain effectively unchanged

subject chooses the strategy that gives them the highest expected utility. The results are given in Table 3.

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 & Weizsäcker, 2008; Polonio & Coricelli, 2019; Rey-Biel, 2009; Sutter et al., 2013), which report consistency levels ranging from 54% to 67%. We add to this literature by showing that, 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, subjects often choose strategies which do not maximise the expected utility of their own payoffs.

All in all, these results suggest that our effort to influence subjects' strategic thinking via increased involvement with the game at hand did not translate to either a high frequency of choices that are part of a Nash equilibrium strategy or high best response rates.

3.2 The role of other-regarding preferences

A possible explanation for the seeming failure to 'think like a game theorist' (Crosson, 2000) 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 e.g., Bolton & Ockenfels, 2000; Charness & Rabin, 2002; Dufwenberg &

Table 5 OLS regression: 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

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

Kirchsteiger, 2004; Falk & Fischbacher, 2006; Fehr & Schmidt, 1999; Rabin, 1993 for social preference models and Cooper & Kagel, 2016; Fehr & Schmidt, 2006; Sobel, 2005 for empirical overviews of the literature). It could be that strategy 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.¹⁰ 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. At the same time, we find that the other player's payoff systematically affects subjects' rankings. In particular, our results reveal that that subjects are, on average, inequity averse (e.g., Bolton & Ockenfels, 2000; Fehr & 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 italicized 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.

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. We find 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

¹⁰ We find a very high level of ranking-consistency across games. To evaluate consistency, we compare the relative ranking of any two own-other payoff pairs that appear in more than one game. For example, the payoff pairs (4, 4) and (6, 6) both appear in games 1, 2, and 3, yielding three pairwise comparisons of their relative ranking across games. Overall, we find that the relative ranking of pairs of payoffs are consistent in 95% of the cases, suggesting that subjects took this task seriously and that they have well-defined preferences over own and other's payoffs in the context of our games.

payoff has a significant positive effect, indicating that, *ceteris paribus*, people dislike inequity.¹¹

To explore whether other-regarding concerns might account for some or many of the departures from own-payoff best responses, we re-calculate optimal strategy 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 (see the All players column of Table A5 in Online Appendix A). On average, however, the best response rate increases only moderately from 54% to 57%, a difference that is nevertheless statistically significant (Wilcoxon Signed-rank test, $p=0.004$; paired t test, $p=0.025$).¹²

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 Iriberry & Rey-Biel, 2013). To test for potential heterogeneity, 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).¹³

¹¹ We obtain the same qualitative results if we use a rank-ordered logistic regression instead. The results are reported in Table A4 in Online Appendix A.

¹² For the test we use an individual's average over all games as the unit of observation.

¹³ Note that different functional forms of other-regarding preferences are also feasible (see e.g., Charness and Rabin (2002) on reciprocity and efficiency concerns, or Engelmann and Strobel (2004) on the role of efficiency and maximin preferences). We opt for inequity aversion because of the patterns we observe at the aggregate level (compare Table 3) and because it has proven to be a strong factor in explaining laboratory behaviour (see e.g., Andreoni and Miller, 2002; Fisman et al., 2007; Iriberry and Rey-Biel, 2013 who estimate a preferences-type distribution and find that inequity-averse individuals are the second most frequent type after the selfish one or Beranek et al., 2015 who find substantial inequity aversion in several samples of students).

On this basis, 55 out of 100 subjects are classified as Selfish, while 44 subjects are classified as Inequity Averse. We find no subject to be Inequity Seeking.¹⁴ This classification allows us to re-calculate the best response rates separately for each type (see the Selfish and Inequity Averse players columns of Table A5 in Online Appendix A). For Selfish individuals we find that, on average, the best response rate amounts to 55%, irrespective of whether using expected payoffs or expected rankings (Wilcoxon Signed-rank 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% (Wilcoxon Signed-rank test, $p=0.004$; paired t test, $p=0.011$). Thus, 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—a reassuring result, as it shows that the ranking task is picking up something that feeds into subjects' strategy choices.

In sum, while we find that taking into account individuals' attitudes to unequal payoffs can indeed improve best response rates, the differences are quite small. The face value interpretation of our data therefore gives only rather limited support for the idea that such preferences are a major explanation for non-optimal behaviour, as conventionally judged in terms of maximising expected own-payoff values. However, some caveats are in order here. In particular, although the ranking task may do a good job of tapping into players' attitudes to inequity, there may be other dimensions of social preferences not captured by this instrument affecting players' strategy choices. For example, revealed preferences in the ranking task might be context dependent: when some pairing of own-other payoffs is being evaluated in the context of where that cell is located in the game matrix, there may be considerations of reciprocity or (un)kindness that are absent from the ranking exercise. Similarly, the order by which players encounter the ranking and the belief-elicitation task might influence their chosen strategies. If the order is reversed, and subjects first consider how specific outcomes depend on their counterparts' intentions, their subsequent ranking may be more aligned with the perceived (un)kindness of these outcomes in the context of the game. Furthermore, some subjects might be prosocial and can reveal this preference in a static scenario, but when the decision becomes strategic, they might believe that their counterparts are selfish, and choose a more selfish strategy instead. We cannot exclude these possibilities, and our data do not allow us to examine any particular consideration of this kind in a systematic manner. Thus, caution should be exercised in extrapolating insights to environments that measure a wider range of social preferences besides inequity concerns.

¹⁴ One subject falls under neither of these categories as this individual 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. To further test the robustness of our results, 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.




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

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 attitudes to inequity. 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 optimal strategies in at least half of the games, only very few people do so in all eight games (see Figure A1 in Online Appendix A for the full distribution). The mean (median) number of optimal strategy choices is 4.33 (4) when using expected payoffs, and 4.59 (4.5) when using expected rankings, a difference that is statistically significant (Wilcoxon Signed-rank test, $p=0.004$; paired t test, $p=0.025$). This confirms that considering subjects' attitudes to inequity increases best response behaviour, but that this effect is only modest in size.

Next, we look at the cost of non-optimal play. We compare the expected payoff (expected ranking) between the chosen strategy and the strategy that would have been optimal given a subject's stated beliefs. Figure 5 shows the percentage of non-optimal strategy 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 of expected payoffs (median £1.8) and 1.4 points in expected rankings (median 1).

Table 6 Regression analysis of optimal strategies

Dependent variable	Optimal strategy payoffs			Optimal strategy 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 strategy		-0.359 *** (0.090)				
Optimal strategy contains lowest payoff (1 if yes, 0 otherwise)			-0.896 *** (0.191)			
Optimal strategy contains highest payoff (1 if yes, 0 otherwise)			0.090 (0.165)			
Std. dev. in expected rankings				0.210 ** (0.095)	0.267 *** (0.099)	0.213 ** (0.106)
Std. dev. in ranks within optimal strategy					-0.368 *** (0.141)	-0.687 *** (0.190)
Optimal strategy contains least preferred option (1 if yes, 0 otherwise)						0.081 (0.214)
Optimal strategy contains most preferred option (1 if yes, 0 otherwise)						
Constant	-0.310 ** (0.151)	0.522 ** (0.252)	0.194 (0.197)	-0.106 (0.150)	0.599 * (0.306)	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	0.016	0.024

This table reports coefficient estimates from logistic regressions. The dependent variable is whether the chosen strategy 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 strategy 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. This finding is further reflected in response times: the bigger the advantage of the best option, the faster subjects reach a decision (see Table A6 in Online Appendix A).

In Model (2), we add the standard deviation of the three possible own earnings within the optimal strategy as a second explanatory variable. The coefficient is significantly negative, indicating that the higher the variation in own payoffs within the strategy 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; Di Guida and Devetag, 2013) showing that choice behaviour is susceptible to the influence of out-of-equilibrium features of the games under consideration. Di 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 strategy 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. This result bears some resemblance with recent evidence by Avoyan and Schotter (2020), who study allocation of attention in experimental games. In line with our results, they find that when presented with a pair of games, subjects pay more attention to the game with the greatest minimum payoff, although in their case the game with the largest maximum payoff also receives increased attention.

In Models (4) to (6), we repeat the same analysis but now using optimal strategies 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 chosen strategies

Our analysis so far has looked at behaviour only in the structured environment, where subjects were encouraged to invest more effort in thinking about the decision situation. Our results have revealed that despite our attempt to induce subjects

Table 7 Comparison of game play based on different models' predictions

Game	Nash			Level-1			Level-2		
	S	U	<i>p</i> value	S	U	<i>p</i> value	S	U	<i>p</i> value
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

Average proportion of chosen strategies 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)

to extensively engage themselves with all relevant elements of the game at hand, strategic sophistication (as measured by the fraction of Nash-consistent chosen strategies) as well as best response behaviour was rather low. To test whether structuring subjects' decision-making process has any influence on strategy choices at all, in the following we compare the patterns of choice in the Structured treatment with those in the Unstructured treatment. The results are reported in Table 7.

Overall, we do not find any 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 subjects play according to the Nash, Level-1, or Level-2 predictions across the two treatments. On average, chosen strategies 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).¹⁵ These results hold if we compare chosen strategies separately for each game and player type (BLUE or RED player) (see Tables B1 and B2 in Online 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.¹⁶

¹⁵ We also do not find any systematic differences when we compare the overall distribution of chosen strategies across treatments for either of the players: out of the sixteen possible comparisons between the Structured and the Unstructured treatment, only one yields a weakly significant result (RED players in game six, $p=0.074$). All other comparisons are not significant at the 10%.

¹⁶ It has been suggested to us that the lack of difference in subjects' behaviour across the two treatments may have been driven by our incentivizing procedure. In the Structured treatment where we employ a single sequence of the eight games using three different tasks, each decision in a game has a probability

4 Concluding remarks

Our study sought to examine the extent to which strategic thinking 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. Although the ranking exercise was not guaranteed to capture all social preferences that might affect strategic choice, it was intended to encourage players to pay attention to all of their own *and* the other player's payoffs, with considerations of the other player's payoffs hopefully leading to more thought being given to the other player's likely choice of strategy. Despite structuring the decision process in this way, a majority of chosen strategies were not consistent with equilibrium behaviour. Furthermore, subjects failed to best respond to their own stated beliefs almost half of the time when judged in terms of maximising own-payoff expected values. Allowing for different levels of risk aversion using a standard utility function form did not help much—although avoiding the worst possible payoff did appear to carry some weight. We further examined the possible role of subjects' attitudes to inequity: we found that making allowance for such preferences increased best response rates, but only marginally.

We went to considerable lengths, within the usual experimental time and budget constraints, to structure the decision process in a way that would facilitate strategic thinking. The fact that, despite such assistance, 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. It might be 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. Kahneman et al. (2016) put this view concisely, when they observed: *'Where there is judgment, there is noise—and usually more of it than you think'*. Thus, some irreducible minimum level of imprecision might generate some proportion of sub-optimal choices. Indeed, 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 apply accumulator mechanisms (e.g., Golman et al., 2020) or preferences for deliberate randomization (e.g. Agranov & Ortoleva, 2017; Cerreia-Vioglio et al., 2019; Dwenger et al., 2018) may also be candidates for further consideration.

Another possibility might be the failure to fully understand the structure of the game and how combinations of strategies produce outcomes (e.g., see Bosch-Rosa & Meissner, 2020; Chou et al., 2009 for similar evidence). Even though the purpose

Footnote 16 (continued)

of 1/24 to be payoff relevant (3 tasks \times 8 games). In the Unstructured treatment, each decision in a game has a probability of 1/96 to be chosen for payment (12 repetitions \times 8 games). Given that subjects were not aware of the total number of games in either of the treatments, we think it is unlikely that the lack of significant differences in subjects' chosen strategies is driven by our incentivizing procedure. However, we cannot rule out this possibility.

of the ranking and the belief-elicitation exercise was to provide subjects with a more thorough rendering of the strategic environment, the manipulation might not have been equally effective for all. A proportion of subjects might have still focused their attention solely in the strategy-choice task due to cognitive inability to process the insights gained from the pre-decision tasks at the same time (see e.g., Bosch-Rosa et al., 2018; Carpenter et al., 2013; Fehr & Huck, 2016 for related evidence on how cognitive abilities matter for strategic sophistication). Such cognitive limitations might explain the use of heuristic approaches that do not comply with the logic of equilibrium behaviour by avoiding for instance the worst possible payoff. Perhaps these possibilities are complementary, and non-equilibrium behaviour is simply too heterogeneous to be captured by a single explanation. We leave this open for future research.

Finally, we want to comment on the lack of any significant difference in the overall patterns of chosen strategies between the Structured and Unstructured treatments. Our finding may reassure researchers that simply asking participants to pick their preferred strategy 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.

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1007/s10683-021-09725-8>.

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. We also thank Kirsten Marx for excellent research assistance.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Agranov, M., & Ortoleva, P. (2017). Stochastic choice and preferences for randomization. *Journal of Political Economy*, 125(1), 40–68.
- Andreoni, J. (1995). Cooperation in public-goods experiments: Kindness or confusion? *The American Economic Review*, 85(4), 891–904.
- Andreoni, J., & Blanchard, E. (2006). Testing subgame perfection apart from fairness in ultimatum games. *Experimental Economics*, 9(4), 307–321.
- Andreoni, J., & Miller, J. (2002). Giving according to GARP: An experimental test of the consistency of preferences for altruism. *Econometrica*, 70(2), 737–753.
- Avoyan, A., & Schotter, A. (2020). Attention in games: An experimental study. *European Economic Review*, 124, 103410.

- Bayer, R. C., & Renou, L. (2016). Logical abilities and behavior in strategic-form games. *Journal of Economic Psychology*, *56*, 39–59.
- Beranek, B., Cubitt, R., & Gächter, S. (2015). Stated and revealed inequity aversion in three subject pools. *Journal of the Economic Science Association*, *1*(1), 43–58.
- 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. (2020). The one player guessing game: A diagnosis on the relationship between equilibrium play, beliefs, and best responses. *Experimental Economics*, *23*(4), 1129–1147.
- Bosch-Rosa, C., Meissner, T., & Bosch-Domènech, A. (2018). Cognitive bubbles. *Experimental Economics*, *21*(1), 132–153.
- 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. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton University Press.
- 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.
- Carpenter, J., Graham, M., & Wolf, J. (2013). Cognitive ability and strategic sophistication. *Games and Economic Behavior*, *80*, 115–130.
- 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.
- Cerreia-Vioglio, S., Dillenberger, D., Ortoleva, P., & Riella, G. (2019). Deliberately stochastic. *American Economic Review*, *109*(7), 2425–2445.
- Charness, G., & Rabin, M. (2002). Understanding social preferences with simple tests. *The Quarterly Journal of Economics*, *117*(3), 817–869.
- 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.
- 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), 35–58.
- Cooper, D. J., & Kagel, J. H. (2016). Other-regarding preferences: A selective survey of experimental results. In J. H. Kagel & A. E. Roth (Eds.), *The handbook of experimental economics* (Vol. 2, pp. 217–289). Princeton University Press.
- 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., & 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. A., & Weizsäcker, G. (2008). Stated beliefs and play in normal-form games. *The Review of Economic Studies*, *75*(3), 729–762.
- Cox, J. C., & James, D. (2012). Clocks and trees: Isomorphic Dutch auctions and centipede games. *Econometrica*, *80*(2), 883–903.
- Croson, R. T. (2000). Thinking like a game theorist: Factors affecting the frequency of equilibrium play. *Journal of Economic Behavior & Organization*, *41*(3), 299–314.
- 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.
- Di Guida, S., & Devetag, G. (2013). Feature-based choice and similarity perception in normal-form games: An experimental study. *Games*, *4*(4), 776–794.
- Dufwenberg, M., & Kirchsteiger, G. (2004). A theory of sequential reciprocity. *Games and Economic Behavior*, *47*(2), 268–298.
- Dwenger, N., Kübler, D., & Weizsäcker, G. (2018). Flipping a coin: Evidence from university applications. *Journal of Public Economics*, *167*, 240–250.
- Engelmann, D., & Strobel, M. (2004). Inequity aversion, efficiency, and maximin preferences in simple distribution experiments. *American Economic Review*, *94*(4), 857–869.
- Falk, A., & Fischbacher, U. (2006). A theory of reciprocity. *Games and Economic Behavior*, *54*(2), 293–315.
- 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.
- Fischbacher, U., Fong, C. M., & Fehr, E. (2009). Fairness, errors and the power of competition. *Journal of Economic Behavior & Organization*, 72(1), 527–545.
- Fisman, R., Kariv, S., & Markovits, D. (2007). Individual preferences for giving. *American Economic Review*, 97(5), 1858–1876.
- 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. B. (2020). The dual accumulator model of strategic deliberation and decision making. *Psychological Review*, 127(4), 477.
- 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*, 79(4), 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). 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.
- Iriberrí, 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.
- Kahneman, D., Rosenfield, A. M., Gandhi, L., & Blaser, T. (2016). Noise: How to overcome the high, hidden cost of inconsistent decision making. <https://hbr.org/2016/10/noise>.
- Krawczyk, M., & Le Lec, F. (2015). Can we neutralize social preference in experimental games? *Journal of Economic Behavior & Organization*, 117, 340–355.
- 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.
- Rabin, M. (1993). Incorporating fairness into game theory and economics. *The American Economic Review*, 83(5), 1281–1302.
- 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., Ortman, A., & Ostadnický, 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.
- Weizsäcker, G. (2003). Ignoring the rationality of others: Evidence from experimental normal-form games. *Games and Economic Behavior*, 44(1), 145–171.
- Zonca, J., Coricelli, G., & Polonio, L. (2020). Gaze patterns disclose the link between cognitive reflection and sophistication in strategic interaction. *Judgment and Decision Making*, 15(2), 230–245.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.