No. 128

THURGAU INSTITUTE OF ECONOMICS
at the University of Konstanz

Irenaeus Wolff

Heuristic Centered-Belief Players

Research Paper Series
Thurgau Institute of Economics
and Department of Economics
at the University of Konstanz

Heuristic Centered-Belief Players[§]

Irenaeus Wolff

Thurgau Institute of Economics (TWI) / University of Konstanz
Hafenstrasse 6, 8280 Kreuzlingen, Switzerland
wolff@twi-kreuzlingen.ch

This version: 30^{th} January, 2023

Abstract

Strategic behavior often diverges from Nash-equilibrium, in particular in unexperienced play. I provide data from a class of simple discoordination games and show that none of the popular models of behavioural game theory predicts the predominant aggregate choice pattern. And yet, Noisy Introspection (Goeree and Holt, 2004) readily accounts for about half of the individual observations. The reason for the apparent paradox and the mismatch of the aggregate data and the models is a disregarded behavioural type that makes up about 25% of the population. These 25% hold beliefs that peak in the centre of the option set and that are roughly symmetric. In addition, the players show a more heuristic process translating their belief into actions, as their choices cannot be explained readily by quantal responding. The behavioural pattern of a 'centered belief' in connection with boundedly-rational decision-making is present also in another prominent game from the literature on behavioural game theory, the 11–20 game. Finally, I show that classifying players as 'heuristic centered-belief types' by one game's beliefs has predictive power for behaviour in the other game.

Keywords: Nash-equilibrium, quantal-response equilibrium, level-k, cognitive-hierarchy, salience theory, noisy introspection, central-tendency bias.

JEL: C72, C92, D83, H41

[§]I am grateful to all the many people who helped me write this paper, in particular Ankush Asri, David Cooper, Miguel Costa-Gomes, Fabian Dvorak, Urs Fischbacher, Moath Hussien, Andrea Isoni, Max Lobeck, Ariel Rubinstein, Julia Werner, the research group at the Thurgau Institute of Economics, participants of the microeconomics seminar at the University of Konstanz, research seminars at Otto-von-Guericke University Magdeburg and Bremen University, as well as participants of FUR 2022 (Ghent), and the 2022 UEA Annual Workshop on Behavioural Game Theory.

1 Introduction

Many interactions in modern societies happen between strangers, and even longer relationships have to start at some point. Thus, understanding interactions that do not have a history is of considerable importance, and for this reason, there is a substantial body of literature studying strategic behavior in one-shot interactions (*e.g.*, Crawford et al., 2013). This research has shown that strategic behavior often diverges from Nash-equilibrium. By now, scholars of behavioral game theory have found an explanation for virtually all instances where behavior diverges from Nash. This paper shows that despite these achievements, none of the popular models is able to predict the predominant aggregate data pattern from a class of simple discoordination games, even though Noisy Introspection (Goeree and Holt, 2004) accounts for the individual behaviour of slightly more than 50% of the population.[1]

The reason for the failure to predict the aggregate pattern is that the models miss an important behavioural type which I call *heuristic centered-belief player*. The behavioural type is characterised by a roughly symmetric belief that puts the highest probability mass on the opponent's 'central' option, as well as a response to their beliefs that clearly diverges from quantal-responding. In the discoordination game, the type breaks the monotonicity predicted by Nash and leads to strong differences in expected payoffs of the available options (in contrast to zero-differences under Nash). However, hardly any participant anticipates these differences correctly, so that only some 10% report a belief for which the best-response coincides with the empirical best-response to the population.

The paper's focal game in its standard form is very simple:

> *Please choose one of the following payoffs, [27, 30, or 33]. You receive the payoff you choose, but only if the other player chooses a different payoff than you. Otherwise, you receive nothing.*

This game is a standard 3×3 normal-form game that mirrors the basic characteristics of economically relevant situations. For example, think of two similar firms which have to decide on which of three different markets they want to enter, or which of three different types of product they want to sell. In addition, suppose they know there will be Bertrand-competition afterwards if they choose the same market, and monopoly rents if they enter different markets.

In the absense of any coordination device, scholars typically look for a symmetric Nash equilibrium as their benchmark. In the case of the discoordination game above, the unique symmetric Nash equilibrium has both players choose 27,

[1] I briefly describe all considered models in Section 4.

30, and 33 with probabilities of 26%, 34%, and 40%, respectively.[2] The frequencies I observe when running the game with 168 participants are 31%, 21%, and 48%.

As pointed out above, deviations from Nash-equilibrium should come as no surprise in such a one-shot environment. In case of such a deviation, however, we should expect the deviation to be captured by at least one of the popular alternative models suggested by behavioural economics. And yet, note that the standard quantal-response-equilibrium, cognitive-hierarchy, level-k, salience-theory, and impulse-balance-equilibrium models also predict monotonic patterns. Team-reasoning brings us back to the symmetric Nash equilibrium, too, because there is no way of distinguishing the players and coordinating players' responses. Interestingly, when I surveyed 80 experimental economists recruited through the Economic Science Association's discussion mailing list, only 14 predicted a monotonically increasing pattern. 16 predicted the type of non-monotonicity I observe, though in sometimes extreme versions, like 50-0-50. 18 predicted a monotonically decreasing pattern, and the mode—32 respondents—predicted an inverted-U pattern.

In this paper, I present the data of the baseline game presented above together with 17 additional treatment variations, to test the robustness of my findings with respect to a whole number of experimental parameters. In all 18 treatments, the second-highest payoff is chosen less often than in the symmetric Nash equilibrium. Moreover, 13 of the 18 experiments yield the particular type of non-monotonic pattern described above, where the modal choice is the highest payoff but the second-highest payoff is chosen less often than the third-highest one. In a 14$^{\text{th}}$ case (a variant with 5 options), it is the third-highest payoff that is chosen less often than the fourth-highest payoff. The remaining 4 treatments show monotonic patterns, and in contrast to the experimental economists' modal prediction, no treatment yields the 'opposite' non-monotonicity of an inverted-U pattern.

A treatment with an incentivised elicitation of participants' reasoning reveals that participants have a "reasonable argument" for over-playing the lowest option in the baseline game. 41% of those who are in favour of choosing the lowest of the three amounts explicitly mention "safety" in justifying their preferred choice. This argument is virtually absent in justifications of other choices, in which only 4% refer to "safety"—but about 20% to "going for the risk", albeit in more varied ways.

[2] There are also two types of asymmetric equilibria, one in which one player chooses "33" and the other chooses "30", and another type in which both players mix, with probabilities $(\frac{10}{11}, 0, \frac{1}{11})$ and $(0, \frac{9}{11}, \frac{2}{11})$. However, as there is no way of breaking the symmetry, I discard them as implausible. Note that no convex combination of the empirical distributions that would result from the different equilibria would yield the non-monotonicity I observe. If at all, we would observe the opposite type of non-monotonicity (where "30" would be the modal choice).

A belief-elicitation treatment sheds more light on the data. While only one third of those choosing either of the lower payoffs play optimally given their elicited belief—compared to 81% of those choosing the highest option—neither group acts in a completely irrational way. If participants deviate from their optimal choice, they almost always deviate to the option with the second-highest expected payoff (in 83% of the cases). And interestingly, this happens to be 27 for 56% of those deviating from 33 (while all of the deviations from 30 go to 33, and 3/4 of those deviating from 27 choose 33).

What this means: even though many people argue as if they were cognitive-hierarchy level-2—“everybody else will be choosing the upper two payoffs”—it’s not a “clean” cognitive-hierarchy reasoning, as the high number of 27s is due to choices that are not best-responses to their beliefs. It is not a pure quantal-response-equilibrium reasoning, either, as that would produce even stronger monotonicity than the symmetric Nash-equilibrium. It is not even non-equilibrium quantal-response behaviour, even though most suboptimal choices are next-best choices: fitting quantal responses given the reported beliefs yields a strongly monotonic prediction, which also is more extreme than the equilibrium pattern.

Noisy introspection (NI)—being a combination of a cognitive-hierarchy-type model with quantal responses—does not provide a satisfying explanation, either.[3] In the calibrated variant of Goeree et al. (2018), the model predicts the non-monotonic pattern for 5 out of the 14 treatments that lead to non-monotonic choice patterns. Maximum-likelihood fits of experiment-wise fitted data-sets produce non-monotonicities for the same number of treatments.

A cluster analysis of the beliefs data shows why. While roughly half of the participants can be classified as having NI-1, NI-2, or NI-3 beliefs, the other half cannot. Among this second half, there is a robustly-classified cluster of about 25% of the whole population who exhibit a form of the central-tendency bias (see, for example, Crosetto et al., 2020, and references cited therein). These ‘heuristic centered-belief players’ show the above-mentioned single-peaked and roughly symmetric beliefs. They choose 27 particularly often, and excluding them from the data would result in a monotonic aggregate data pattern.

Heuristic centered-belief players are present also in a data set on the 11–20 game provided by Goeree et al. (2018), but only when the options follow their natural order.[4] There are no centered-belief players in treatments in which the order of options is perturbed (19–18–17–...—12–11–20 and 14–13–12–11–19–18–17–16–15–20). Again, the centered-belief players’ behaviour differs markedly from other participants’ behaviour, accounting for 50-67% of all—non-optimal—

[3]Combining quantal responses with the cognitive-hierarchy model rather than with level-k (as in noisy introspection) fits the data worse.

[4]In the 11–20 game, players obtain the amount they request between 11 and 20 currency units, plus a bonus of 20 units in case they request exactly 1 unit less than their opponent.

choices below 17, even though they make up for only 8-17% of the population.

In a final treatment variation, I test whether heuristic centered-belief play is a strategy that people sometimes adopt when playing games like the ones studied here or whether heuristic centered-belief players are a type of person who routinely use such a strategy accross situations. In particular, I let participants play a 17–20–23 game, ask for their beliefs in that game, and then let them play the 11–20 game. I run the same cluster analysis as before on the 17–20–23 beliefs, categorising participants into centered-belief players and other players, and then look at the two groups' behaviour in the 11–20 game. The centered-belief players again are much more likely to choose an action below 17 compared to all other players, in this case, by a factor of three. Hence, I again find a fraction of roughly 25% heuristic centered-belief players (who produce a non-monotonic pattern, choosing "17", "20", and "23" with relative frequencies of 28%, 12%, and 60%, respectively). And second, heuristic centered-belief play seems to be a strategy that a certain part of the population applies across games.

Looking at the broader picture, this paper does a number of things: it establishes the existence of a neglected behavioural type. It shows that the type exists not only in extremely contrieved games but in very standard games like discoordination games and the 11–20 game. Also, the type is unlikely to be a spurious finding, as it exists in several different games, and as people showing the type behave in a distinct way both when reporting beliefs and when taking actions, which on top can be documented across games.

More generally, the paper suggests that the literature typically has focused too much on 'homogeneous' models which may contain different types, but where, by and large, all types follow the same type of logic (as, for example, in cognitive-hierarchy-type models). Future research still will need to specify the exact conditions of when the type will show and what the exact mechanism is behind their choices. However, my paper suggests that both explaining and predicting behaviour in related games needs to account for a fraction of heuristic centered-belief players.

2 Related literature

This paper contributes to three strands of literature. First, there is a large and growing literature on strategic reasoning in one-shot games (for an introduction, *cf.*, Crawford et al., 2013). In most of this literature, there is a particular model of behavioural game theory that explains behaviour best, but the best-performing model differs between papers (and sometimes, even within papers, as in Bardsley et al., 2010). In this sense, my paper falls into the same category as, for example, Goeree et al. (2018), showing that noisy introspection accounts for the behaviour of a substantial part (one half) of the population particularly well.

However, my paper also shows that participants' aggregate response pattern (and a non-negligible sub-population's individual response patterns) cannot be explained readily by any of the popular models of behavioural game theory. Instead, behaviour is governed by a more heuristic decision-making procedure. In this sense, the paper is related to papers like Apesteguia et al. (2007) or Eliaz and Rubinstein (2011), who show that in certain situations, simple decision rules may explain behaviour best. Within this literature strand, closely-related papers are Crosetto et al. (2020), Wolff (2021), and Sontuoso and Bhatia (2021).

Crosetto et al. (2020) have their participants bid for an object with known value in a two-player auction against a computer player whose bid is determined by uniform randomisation. After that, participants have to indicate the probabilities with which the computer's bid falls into one of five evenly-sized bins of possible bids. Even though the computer's uniform-randomisation strategy is known to participants, a majority reports "beliefs that have a peak in the interior of the range". My study adds to the findings of Crosetto et al. (2020) in several ways. I distinguish 'centered-belief players' from other participants whose beliefs have a peak in the interior but are clearly asymmetric. I further document some limits to the phenomenon, as it does not seem to be present when actions do not appear in their natural order. Most importantly, I relate participants' beliefs to their actions.[5] In particular, I show that 'centered-belief players' differ qualitatively in their behaviour from the rest of the population, that this differing behaviour affects the overall data pattern in two different games, and that the types' strategies in the two games are correlated within-participants.

Wolff (2021) studies symmetric pure discoordination games in which participants' pure strategies can be distinguished only by their labels and positions. The study finds that participants play the pure discoordination games as if they were giving up on reasoning strategically and betting on one of the strategies.[6] Similarly, Sontuoso and Bhatia (2021) study coordination games and hide-and-seek games in which options are denoted by natural-language words. They show that players often choose actions with a frequently-used word as a label when they have an incentive to match their opponent's action. In contrast, players who have to mismatch their opponent's action rely less on actions that have a frequently-used word as a label.[7] I explicitly vary the salience of the options be-

[5] Crosetto et al. (2020) report only that actions and beliefs are correlated between participants.

[6] Note that—unsurprisingly—this explanation fails to account for the data presented in this paper (Wolff, 2021, proposes that agents resort to betting-like behaviour only in case strategic thinking leads to the conclusion that any option is as good as any other, which will not be the case here). If participants have to bet on "27 €", "30€", or "33€"to receive a prize of 10 Euros in case a random draw selects the same option, 18% bet on 27, 38% bet on 30, and 44% bet on 33.

[7] See, *e.g.*, Mehta et al. (1994), Bardsley et al. (2010), Faillo et al. (2017), or van Elten and

tween treatments, without any apparent effects on the resulting behaviour. In that sense, the monetary incentives seem to be strong enough to dominate any 'prominent-number effects'.

The second literature strand looks at the stability of behaviour across different games, as, for example, in Georganas et al. (2015) or Rubinstein (2016). Georganas et al. (2015) show that participants exhibit little stability in terms of their estimated level-k types across two families of games.[8] In contrast, Rubinstein (2016) finds positive correlations between a whole range of games in terms of 'decision styles': the fraction of a participant's "contemplative" choices in collections of nine games in most cases is predictive of the same participant's choice in a tenth game. Similarly, my findings also suggest that decision strategies may not be that different between games. Identifying 'centered-belief players' in a setting where players want to coordinate on different actions is predictive for choices in a setting where players benefit hugely from outsmarting the other player (and where there always is at least one player who would want to deviate from her action *ex-post*).

Last but not least, the paper contributes to the literature studying competitive situations in which collusion would be beneficial. This literature centers around settings in which firms are active in the same market and studies topics such as communication, the number of firms, or the particular market mechanism (Haan et al., 2009). My paper studies a stylised setting in which firms want to divide a market amongst themselves by offering diversified products or several markets when these markets yield unequal monopoly rents. Thus, rather than looking at a homogenous good and price or quantity choices, one could think about the setting in this paper as being equivalent to firms choosing to market homogenous goods in different markets or endogenously-heterogeneous goods with a known Bertrand pricing strategy afterwards.

3 Experimental design and setups

As stated in the introduction, I ran 18 treatment variations on the basic discoordination game. In the game, players obtain the amount they request as long as their opponent requests a different amount. The initial focus of the study was a comparison of the baseline treatment with the game by Berger et al. (2016). In their game, there are many players who have to choose one of three available payout amounts. The amount chosen by the lowest number of players is the "winning amount", and one player would be drawn for payment from the group of players who chose the winning amount.

Penczynski (2020) for studies on strategic reasoning in coordination games.

[8]See Cooper et al. (2018) and Hyndman et al. (2022) for similar results.

Because the project started as an exploratory exercise, the initial treatments were added to unrelated experiments. Having seen the surprising data pattern in the baseline treatment, I started exploring the finding's robustness in a variety of treatments. To keep things comparable, I kept running the treatments as an extra task in unrelated experiments for most treatments. To test whether this arrangement was driving the results, I also ran treatments as the first or even the only part of experimental sessions. This was the case in 5 of the 18 data sets of Table 1. In addition, the game was the second main part of the experiment in the team version with options 40-44-48 (out of two, after an unrelated one-shot public-good task played with a different opponent).

The first round of robustness checks focused on the numbers themselves. These treatments would not have "30" as the middle option, or increase the differences between payoffs. Then, I ran two treatments where all pairs of participants would be paid rather than only one randomly selected couple. Paying all participants meant reducing the prizes for budget reasons. This was followed by a number of treatments with more options, and online replications conducted on Ariel Rubinstein's website http://gametheory.tau.ac.il/ using hypothetic payoffs, very large payoff differences, and (mostly) former students from his game-theory courses.

As part of another study (Bauer and Wolff, 2018), I also ran sessions on two variants of the game in which participants would play the same game simultaneously with all other participants in their session, with the restriction that they had to choose the same action in all the games. The difference between the variants was how the game was presented. In an 'opponent frame', the setup was presented as a two-player game first, followed by the explanation that participants would play the two-player game against all others in the session simultaneously (being paid their average payoff). In a 'population frame', in contrast, they were told that they would be paid their chosen amount for every other participant in the session who chose a different amount, divided by the number of participants in the session.

I further conducted an experiment using 2-player teams as a means to incentivise the disclosure of participants' reasoning. The approach, pioneered by Burchardi and Penczynski (2014), rests on the idea of letting one team member suggest an action for the game and write a justifying message to the other member who then takes the final decision for the whole team. To economise on research money and participants, it is not revealed who in the team is the suggesting player and who is the player making the final decision, so that both players suggest an action and write a message. Importantly, though, the communication along the actual path of play remains a one-way communication.

This setup yields two types of choices, suggestions and implemented actions after participants have received the suggestions and messages of their fellow team member. In terms of choice data, I will focus on the suggestions as they are uninfluenced by others. Focusing on implemented choices instead would boost the non-monotonicity of the observed pattern of choices strongly.

Finally, let me describe the design features of the treatments in which I elicited beliefs. First, I added a second additional part to two of the initial sessions. In the sessions, I asked 58 participants of a 27–30–33 treatment to report their estimate of the probabilities with which the other participants would choose each of the three options. They would be paid an additional 2 Euros if their estimates of others' behaviour did not differ for any of the options by more than 2 percentage points.

The final four sessions I ran had two main parts and the belief-reports part in between. In the first part, the 99 participants would play a 17–20–23 treatment (in points, with a conversion rate of 4 points per Euro). Then, they would go through the belief elicitation described above, and finally, they would play the standard two-player 11–20 game. In the game, each player would obtain the amount they request unless they ask for exactly 1 point less than their opponent. If they choose exactly 1 point less than their opponent, they receive an additional bonus of 20 points. One of the two games would be selected randomly for payment. The final payoff would consist of the payoff from the selected game and the belief elicitation, plus a fixed participation fee of 2.50 EUR.

With the exception of the three hypothetic-incentives treatments run on http://gametheory.tau.ac.il/, all sessions were run with the local participant pool of the LakeLab at the University of Konstanz. I used z-Tree (Fischbacher, 2007) to programme the experiments, and hroot (Bock et al., 2014) for recruitment. For the online sessions of the 9–10–11, 17–20–23, and 40–44–48 treatments, I additionally had to use z-Tree-unleashed (Duch et al., 2020). Table 1 presents an overview of all experiments, their particularities, and the data.

4 Predictions

For the predictions, I use the symmetric Nash-equilibrium and the popular models from behavioural game theory. The standard (logit) quantal-response-equilibrium, cognitive-hierarchy, level-k, salience-theory, and impulse-balance-equilibrium models all predict monotonically-increasing patterns. Team-reasoning brings us back to the symmetric Nash equilibrium, too, because there is no way of distinguishing the players and coordinating players' responses. Only noisy introspection can account for non-monotonic patterns in some cases.

Options	Option 1	Option 2	Option 3	Option 4	Option 5	N.obs.	Lab/Online	Reason/Particularities	Payment
27-30-33	34.5	**20.9**	44.5			110	lab	Baseline treatment	pay one pair
27-30-33	24.1	**20.7**	55.2			58	lab	belief elicitation (+ replication)	pay one pair
24-27-30	34.6	**15.4**	50.0			52	lab	not having "30" in the middle	pay one pair
24-30-36	29.8	**23.8**	46.4			84	lab	increasing the payoff differences	pay one pair
20-30-40	19.3	**18.1**	62.7			83	lab	increasing the payoff differences further	pay one pair
5-6-7	17.5	30.0	52.5			40	lab	paying all	pay all
7-7.5-8	21.2	**10.6**	68.2			66	lab	paying all	pay all
9-10-11*	17.6	29.4	52.9			85	online	benchmark for survey	pay all (random part)
17-20-23*,‡	24.2	**21.2**	54.5			99	online	predictive power of 'centered-belief types'	pay all (random part)
6.7-7.2-7.7-8.2	6.9	34.7	**13.9**	44.4		72	lab	increasing the number of alternatives	pay all
5-5.5-6-6.5-7	2.1	12.7	17.6	21.8	45.8	142	lab	increasing the number of alternatives	pay all
6.2-6.7-7.2-7.7-8.2	2.3	9.1	17.0	**15.9**	55.7	88	lab	...without having 'round' payments	pay all
5.4 − 6.3 − 7.2*	25.5	**25.3**	49.2			308	online	'replication' amongst GT students	no payment
6000-7000-8000*	27.5	30.0	42.5			209	online	...with large differences	no payment
5000-6000-7000-8000-9000*	8.1	16.1	**14.4**	25.5	36.5	236	online	large differences, 5 options	no payment
27-30-33	38.0	**24.1**	38.0			108	lab	discoordination with everybody, 'population frame'	pay one pair
27-30-33	29.2	**14.2**	56.6			106	lab	discoordination with everybody, 'opponent frame'	pay one pair
40-44-48	28.9 (25.0)	**26.3** **(15.8)**	44.7 (59.2)			76	online	discoordination in teams, suggestion (decision)	pay all (random part)
27-30-33	18.4	34.7	46.9			98	lab	'replication' of the Berger et al. game	pay the winner

Table 1: Overview of the data. Non-monotonicities are marked in bold-face. *The game was the first or only task of the experiment. ‡There was a conversion rate of 4 experimental tokens to 1 Euro.

Nash-equilibrium. I focus on the symmetric equilibrium (for the focal game, 79/299, 101/299, 119/299) because there is no way of breaking the symmetry for the players. Just to name them, the asymmetric equilibria are as follows: (a) one player claims the highest payoff, the other player claims the second-highest payoff, or (b) one player claims 33 with 2/11 and 30 with 9/11 probability, while the other player claims 33 with probability 1/11 and 27 with probability 10/11.[9]

(Logit) Quantal-Response Equilibrium. Under a quantal-response equilibrium, agents make errors, but the likelihood of the errors decreases in their costs. To model such behaviour, agents typically are assumed to follow a logistic-choice function ("quantal response") in their decisions. The probability of choosing an action a_j is thus given by:

$$\Pr(a_j) = \frac{e^{\lambda \pi_j(\mathbf{s_{-i}})}}{\sum_k e^{\lambda \pi_k(\mathbf{s_{-i}})}},$$

where $\pi_j(\mathbf{s_{-i}})$ is the (expected) payoff of choosing a_j when all other players choose according to $\mathbf{s_{-i}}$. The equilibrium then assumes that players mutually take their noisy choice behaviour into account, so that each player plays a quantal response to the quantal response(s) of the other player(s). The model has a single parameter, λ, which measures the degree of rationality (where $\lambda = 0$ means uniform randomization and $\lambda \to \infty$ means perfect best-responding). Figure 1 displays the model's predictions for the baseline game depending on λ. As becomes clear from the figure, the QRE comes closest to a non-monotonic pattern when $\lambda = 0$ and all players randomize uniformly.

Level-k. Under a level-k model, participants have different types, so-called "levels of reasoning". Level-0 types are assumed to react in an intuitive manner to their set of strategies, not taking into account anything that concerns the other player. This may mean that they either randomize uniformly over their set of strategies, that they pick a particularly salient strategy (such as "YES!!" in a set that consists of "YES!!", "maybe", "perhaps", and "who knows"), or the strategy that intuitively suggests the highest payoff (as in the 11–20 game by Arad and Rubinstein, 2012). Levels $k > 0$ then always play a best-response to the corresponding level $(k - 1)$, so that level 1 best-responds to level 0, level 2 best-responds to level 1, and so on. Proponents of level-k commonly assume that level 0 only exists in agents' minds. Similarly, very high levels ($k > 4$ or even $k > 3$) are usually excluded because they seem implausible, and proponents of the level-k model expect the level-distribution to be hump-shaped.

[9] See ftn. 2.

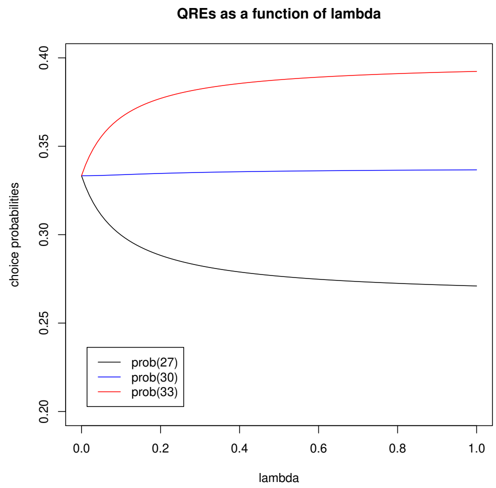

Figure 1: Logit Quantal Response Equilibria as a function of λ. For λ-values greater than 1, the QRE barely changes and approaches the symmetric Nash-equilibrium monotonically.

For our game, there are two 'standard' ways of modelling level 0. Either, we assume level 0 to randomize uniformly, or we assume level 0 to pick "33" because it is particularly salient and at the same time the option that intuitively suggests the highest payoff like in the 11-20 game.[10] If level-0 randomizes uniformly, uneven levels choose "33" and even levels choose "30". If level-0 chooses "33", uneven levels choose "30" and even levels choose "33". In either case, "27" is never played (or at most by 1/3 of the "non-existent" level-0 types). And in neither case should the response distribution display a U-shape.

Cognitive Hierarchy. The cognitive-hierarchy model (Camerer et al., 2004) is closely related to the level-k model. The difference is that levels $k > 0$ believe that the population of other players consist of levels 0 to $(k-1)$ (rather than $[k-1]$ alone), and that this belief follows a truncated Poisson-distribution. To predict behaviour in new games, the Poisson parameter has been suggested to be $\tau = 1.5$. Figure 2 displays the predictions as a function of τ. What Figure 2 shows is that the response distribution should not have a U-shape. The only non-monotonicity that could be generated (for a $\tau \approx 1.1$ and a "33"-choosing level 0) would be a very slight hump-shape.

[10] We could add a variant under which "30" might be salient because it's in the middle and a round number. Doing so would not change the results in any meaningful way.

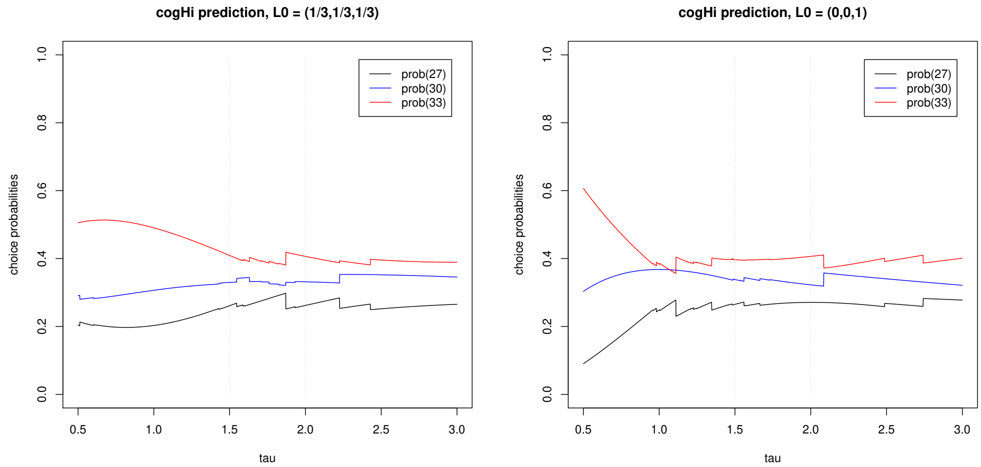

Figure 2: Cognitive-hierarchy predictions as a function of τ. $\tau = 1.5$ has been suggested as a good predictor for new games, and $\tau > 2$ tend to be considered as implausible.

Salience theory. Salience theory (Bordalo et al., 2012) was conceived originally for choice under risk. However, given its growing popularity, I also include it here by calculating a standard Nash-equilibrium for players who act according to salience theory. The central idea is that agents play a best-response to a distorted version of reality. In particular, it is the probability weights that are distorted by salience. A *state* (that is, an action of the other player) is salient *for the player's own action* when the payoff in that state-action combination deviates strongly from the average payoff of that state:

$$\sigma\big(u(a_j, a_i), \overline{u}_j\big) = \frac{|u(a_j, a_i) - \overline{u}_j|}{|u(a_j, a_i)| + |\overline{u}_j| + \varepsilon},$$

where $u(a_j, a_i)$ is the player's payoff if the opponent chooses his pure strategy a_j and the player chooses a_i, $\overline{u}_j$ is the average payoff if the opponent chooses a_j and the player mixes uniformly over her action set, and ε is a small constant that prevents weird predictions in case $|u(a_j, a_i)| + |\overline{u}_j|$ happens to be close to 0.

The salience weights σ within each action a_i (that is, across the opponent's action set) are converted into ranks r_{ij} (starting at 0 for the most salient state) which in turn determine the decision weights ω_{ij}:

$$\omega_{ij} := \frac{\delta^{r_{ij}} \Pr(a_j)}{\sum_{k=1}^{K} \delta^{r_{ik}} \Pr(a_k)},$$

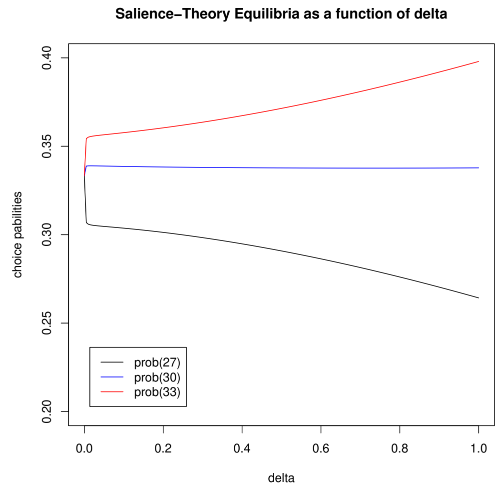

Figure 3: Salience-theory predictions as a function of δ. δ has been estimated to be around 0.7.

where δ is a rationality parameter that has been estimated at ≈ 0.7. Figure 3 shows the equilibria that result under a salience-theory equilibrium as a function of δ. Again, the figure shows that also salience theory predicts a data pattern that monotonically increases in the claimed amount.

Impulse-balance equilibrium. Impulse-balance equilibrium is a concept that is applicable only to settings where actions have a natural order (see, for example, Selten and Chmura, 2008). The discoordination games I am studying have such an order, so that impulse-balance equilibrium is applicable here. The concept is based on learning direction theory.

In learning direction theory, agents compare the payoff they got from their chosen action to the payoffs they might have got under different actions. Then, agents adjust in the direction of higher payoffs. Impulse-balance equilibrium uses the idea, proposing that each alternative action 'emits' impulses. These impulses are influenced by the others' strategies. An impulse-balance equilibrium is reached when for each player, upward- and downward-impulses exactly balance out, given the strategies of their opponents.

Chmura et al. (2014) generalise impulse-balance equilibrium to 3×3 games and apply it to a "Bailiff and Poacher Game" in which the "Poacher" wants to steal fish from one of three ponds. The fish may be differently valuable, and getting caught at either pond is equally bad. Hence, the "Poacher" faces the same payoff structure as players in the discoordination games I study. Thus, I can use the Poacher's impulse-proportionality equation (equation 5 in Chmura et al.) directly:

$$q_i = \frac{(1-p_i)V_i}{(1-p_1)V_1 + (1-p_2)V_2 + (1-p_3)V_3}, \quad \text{for } i = 1, 2, 3,$$

where $V_1 = 27, V_2 = 30, V_3 = 33, p_i$ is the probability with which the respective actions are chosen by the other player, and q_i is the focal player's probability of choosing the i'th element of $\{27, 30, 33\}$. Applying symmetry, $q_i = p_i$. The resulting system of equations can be solved easily and yields, for example, $(p_1, p_2, p_3) \approx (0.31, 0.33, 0.36)$ for the 27–30–33 treatment.

Noisy introspection. Noisy introspection combines quantal-responding with Level-k, fixing level-0 at uniform randomization (Goeree and Holt, 2004). In particular, *NI*-0 mixes uniformly, and any *NI*-k with $k > 0$ plays a quantal response to *NI*-$(k-1)$. Following Goeree et al. (2018), levels are assumed to follow a Poisson-distribution (truncated at the ninth level) with parameter τ. In principle, the model is able to accommodate the type of non-monotonicity I observe. Using the parameter estimates in Goeree et al. (2018) for prediction, the model predicts a non-monotonic pattern for five out of the 14 non-monotonic treatments (for the four 27–30–33 treatments, the prediction is relatively far off, with $(p_1, p_2, p_3) \approx (0.29, 0.28, 0.43)$, while for the 40–44–48 treatment, it is almost perfect at $(p_1, p_2, p_3) \approx (0.29, 0.26, 0.46)$). Fitting the data, this number does not increase: again, five out of the 14 non-monotonic treatments yield a non-monotonic fit.[11]

5 Results

I start by focusing on the three-option games first, where I observe a U-shape in 11 out of 14 data sets. Abstracting from noisy introspection, the closest probability distribution to a U-shape that (at least) some of the models accommodate as a special case is uniform randomization. Under uniform randomization, the likelihood of observing at least 11 out of 14 U-shaped data sets given the sample sizes I used is 0.2%. In other words, I can reject even uniform randomization at the standard significance levels.

To assess whether noisy introspection is able to fit the data sufficiently well, I generate 100'000 random data sets using the choice probabilities of the fitted

[11]The five treatments are two of the 27–30–33 treatments with $(p_1, p_2, p_3) \approx (0.29, 0.27, 0.45)$ and $(p_1, p_2, p_3) \approx (0.27, 0.26, 0.47)$, respectively; the 7–7.5–8 treatment with $(p_1, p_2, p_3) \approx (0.27, 0.23, 0.50)$; the 6.7–7.2–7.7–8.2 treatment with $(p_1, p_2, p_3, p_4) \approx (0.09, 0.25, 0.19, 0.47)$; and the 40–44–48 treatment with $(p_1, p_2, p_3) \approx (0.29, 0.26, 0.45)$.

noisy introspection model and calculate the mean squared deviation of the random data sets and the fitted probabilities. Next, I calculate the mean squared deviations of the actual data sets from the fitted probabilities. Even when I fit the noisy-introspection model on each treatment individually, only slightly more than 1% of the 100'000 iterations have a mean squared deviation that is at least as large as the actual data set. Imposing more structure (such as assuming a common Poisson parameter across treatments) reduces this simulated p-value even further. In other words, also the noisy introspection model cannot accommodate aggregate behaviour. Having looked at the general data pattern across all (three-option) treatments, let me explore the treatments in a more fine-grained way next.

5.1 Robustness checks

The first check was to remove "30" from the center. Changing the options to 24–27–30 (slightly) increased the non-monotonicity rather than reducing it. Second, I increased the differences from 27–30–33 over 24–30–36 to 20–30–40. While there is little difference between the 27–30–33 and the 24–30–36 data, the non-monotonicity becomes weaker in the 20–30–40 treatment. However, it does not do so in the expected way: increasing the differences does not increase the prevalence of "30"-choices in the direction of the theoretic predictions. Instead, there is a shift away from the lowest option that in the most extreme (20–30–40) treatment goes entirely to the highest option.

Paying all (for sure) eliminates the non-monotonicity in one of the instances (5–6–7) but not the other (7–7.5–8). Paying all in case the part is randomly drawn for payment acts similarly: it eliminates the non-monotonicity in the 9–10–11 case but does not do so in the 17–20–23 case nor the team treatment (40–44–48).

Introducing additional options leads to a similar picture. In two out of three lab treatments, the distribution exhibits a dip at the second-highest option, and in the third treatment, the highest available option still is chosen far too often compared to the standard equilibrium prediction.

Using a larger sample of former students of Ariel Rubinstein's game-theory courses moves the data closer to a monotonic pattern, too. However, in two out of three treatments, the data still exhibit a non-monotonicity in terms of a dip, even though in one of them, the dip occurs at the third-highest option.

Finally, the non-monotonicity is not a consequence of conducting the treatments as supplementary parts to other experiments: among the five treatments that had the game as the first or only part of the experiment, again three led to a non-monotonicity.

Summing up, none of the variations was able to eliminate the non-monotonicity reliably. I observe a non-monotonicity in 14 out of 18 treatments, and compared

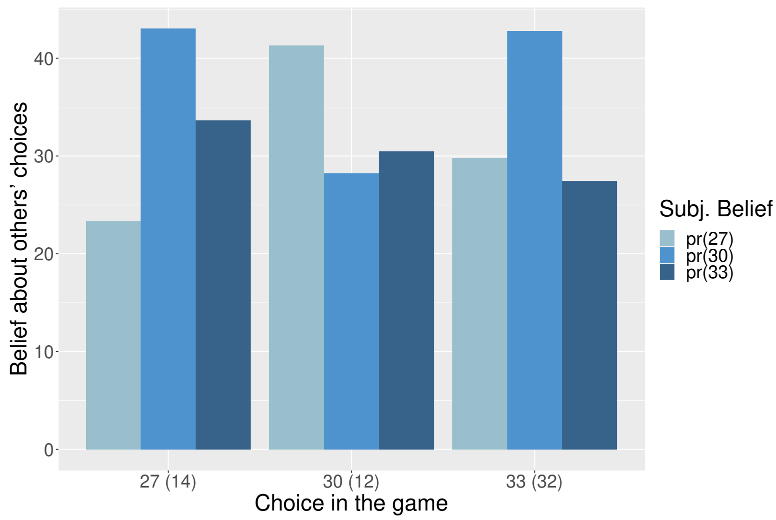

Figure 4: Average beliefs, by participants' own choice (numbers of observations in parentheses).

	Actual choice, in (row-wise) percentages		
Best response (N. obs.)	"27"	"30"	"33"
"27" (9)	55.6	11.1	33.3
"30" (7)	0	57.1	42.9
"33" (42)	21.4	16.7	61.9

Table 2: Participants' choice, by the best response to their reported belief.

to the symmetric Nash equilibrium, the second-highest option is 'under-played' in *all* 18 treatments. Similar findings hold for the other models except for noisy introspection; noisy introspection is able to fit a non-monotonic pattern in 5 out of the 18 treatments. These observations clearly speak in favour of taking the general phenomenon seriously. To understand behaviour in games, it often is very helpful to look at beliefs on top of actions. This is what the next two Sections are focused on.

5.2 Best-responses (based on participants' belief reports)

From 58 participants of a 27–30–33 treatment, I also elicited the beliefs regarding the other participants' choices afterwards.[12] Figure 4 shows the average beliefs, depending on participants' choices. Contrary to the actual choice pattern, the majority of participants seem to agree that most others will choose "30" ("30" carries the largest probability mass in the average belief of the '27–choosers' as well as that of the '33-choosers'). If we now look at what participants should have chosen given their belief, we obtain Table 2. As we can see on the diagonal, 56-62% of the participants choose a best-response to their belief in terms of expected payoffs, irrespective of what that best-response might be.

Looking at choices that were not best-responses, roughly 30% play their second-best choice (43% for those whose best-response is "30"), and 10% play their 'worst-response' (the figures cannot be read from Table 2). These figures suggest that, by and large, participants are acting in agreement with their beliefs, subjects to decision-errors of a quantal-response/logit-choice type. This suggests that the reason for QRE's failure to predict the data lies in the equilibrium assumption rather than in the way beliefs translate into actions. And conversely, quantal-response behaviour seems to be an integral part of the explanation for participants' choices. So far, we have looked at individual belief reports. The next session looks at whether there are special types of participants in terms of their beliefs, to see whether I can relate the belief types to the types from any of the game theoretic models I considered in Section 4.

5.3 Heuristic Centered-Belief Players

The cluster analysis of participants' beliefs reported in Table 3 yields additional insights. First of all, adding up the first three lines (for NI-1, NI-2, and NI-3) suggests that slightly more than half of the participants' beliefs correspond to a type from the noisy-introspection model. Second, looking at lines 4, 5, and 7, just short of 40% believe that "30" is the choice most others will choose (and about three quarters if we take the exact figures in row 1 by their face value). More importantly, though, almost a quarter of the population can be classified as having a 'centered belief' - a belief that places the highest probability on the central option and is (roughly) symmetric around it. And looking at the distribution of actual choices in columns 5–7, it is the 'centered-belief players' who are driving the non-monotonicity of the overall pattern (the only other group exhibiting a

[12] They would be paid an additional 2 Euros if their estimates of others' behaviour did not differ for any of the options by more than 2 percentage points.

Belief on...			Frequency	Actual choice			
"27"	"30"	"33"	(in %)	"27"	"30"	"33"	classification
32.0	35.7	32.3	34.5	4	3	<u>13</u>	$\approx$ NI-1
18.4	33.7	48.0	13.8	<u>5</u>	2	1	NI-2
41.5	14.8	43.7	5.2	0	<u>2</u>	1	NI-3
22.0	52.0	26.0	22.4	5	1	<u>7</u>	'centered belief'
36.7	46.1	17.2	12.1	0	1	<u>6</u>	?
52.2	24.7	23.1	6.9	0	2	<u>2</u>	?
7.5	92.5	0	3.4	0	0	<u>2</u>	outlier
98	0	2	1.7	0	1	<u>0</u>	outlier

Table 3: Results of a cluster analysis of participants' beliefs, alongside the corresponding choice frequencies (best-responses are underlined).

slight dip are the NI-1s).[13]

To test how relevant the result of a meaningful population of centered-belief players is, I looked for another data set of a game with a limited number of options that would have a natural order. The very obvious candidate was Arad and Rubinstein's (2012) 11–20 game in which each player would receive the amount (s)he requests, plus a bonus of 20 in case the other player chooses exactly one monetary unit more. Goeree et al. (2018) provide a data set that includes elicited beliefs for the third game that participants play. In the following Section, I run a cluster analysis on their data.

5.4 Heuristic centered-belief players in other data sets

Running a cluster analysis on Goeree et al.'s (2018) data from the standard version of the 11–20 game (determining the number of clusters again by the SD index) yields 10 clusters, including a 'centered-belief' one. The cluster comprises 8% of the population; however, this number doubles if we use, for example, 6 clusters as suggested by a number of other indices. Notably, this cluster appears only in the variant in which the options follow their natural order, but not in two

[13]Note that 8 clusters is the optimal number of clusters according to SD index (which corresponds to a weighted sum of the average "point scattering" within clusters and the inverse of the total seperation between clusters). Other indices suggest different optimal cluster numbers. However, the 'centered-belief' cluster is the only robust cluster that is always present (tested for a single cluster all the way to 10 clusters). In addition, the robustness check shows that the 22.4% reported in Table 3 are a *lower* bound. Table A.6 in the Appendix shows the proportions of participants whose beliefs are classified as 'centered beliefs' for different numbers of clusters.

	"11"	"13"	"15"	"16"	"17"	"18"	"19"	"20"
Centered-belief players ($n = 6$)	0	0	67	33	0	0	0	0
Other types ($n = 66$)	2	2	3	3	23	36	23	9

Table 4: Choice percentages in the standard (round-3) 11–20 game by belief cluster (beliefs were elicited only for round 3).

treatments in which the order is perturbed (19–18–17–...–12–11–20 and 14–13–12–11–19–18–17–16–15–20). This means that options being naturally ordered seems to be a pre-condition for participants to come up with 'centered beliefs'.

Again, the centered-belief players exhibit a rather peculiar choice pattern, as can be seen from comparing the first row of Table 4 with all other rows. Strikingly, all choices from centered-belief players are below 17 (representing 50% of all such choices; recall that they make up for only 8% of the population; using the above-mentioned 6 clusters, the numbers would go up to 17% of the population making up for 2/3 of the choices below 17). In contrast, slightly less than 10% of the choices of other players are below 17. [14] Note that these choices are clearly dominated: judging by the (clustered) belief, the optimal choice of 19 yields a payoff (20.5) that surpasses the most common choice in that group, 15 (payoff: 18.1), by almost 2.5, whereas the numbers 17 and above all yield payoffs of at least 19.1. This means that in the 11–20 game, centered-belief players clearly do not follow the general quantal-response pattern.

Having found the centered-belief type in two different games, I set out to test the robustness of the findings in a new experiment, and whether classifying participants as centered-belief players in one game is predictive of their choices in the other. The following Section shows that this indeed is the case.

5.5 Centered-belief play as a general strategy

To analyse the out-of-game predictive power of the centered-belief-player categorisation, I ran four sessions with two parts. The first part is a 17–20–23 treatment with belief-elicitation and the second part is the 11–20 game.[15] The hypothesis to be tested was that participants classified as centered-belief players

[14] A Boschloo test on whether choices below 17 are equally prevalent among centered-belief players and other types yields $p < 0.001$. Focusing on choices only from the first round of play in their setup, the 8% centered-belief players still account for one third of the choices below 17.

[15] The sessions lasted for only about 20 minutes, and so I used an exchange rate of 4 points per Euro. If the first part was selected randomly, participants would get a fixed fee of 2.50 Euros on top of their earnings.

	"13"	"14"	"15"	"16"	"17"	"18"	"19"	"20"
Centered-belief players ($n = 24$)	4	8	0	17	8	46	13	4
Other types ($n = 74$)	0	1	1	7	18	49	16	8

Table 5: Choice percentages in the standard 11–20 game, by belief cluster from the 17–20–23 treatment.

by their beliefs in the first game would again make up for a disproportionally large share of the choices below 17 in the 11–20 game. Recall that in Goeree et al.'s data, '11–20-game centered-belief players' made up for 8%(-17%) of the population but for 50%(-67%) of the choices below 17.

Table 5 presents the results for the across-game analysis. While the effect is weaker, it clearly remains present: centered-belief players are more than 3 times as likely to choose a number below 7 compared to all other participants (29% as opposed to 9%), again making up for 50% of all such choices.[16] The findings show two things. First, I again find a substantial fraction of centered-belief players (that is very similar in size to what I observed among the 58 earlier participants) who produce a non-monotonic pattern, choosing "17", "20", and "23" with relative frequencies of 28%, 12%, and 60%, respectively. And second, in the sense of the tested hypothesis, reporting a 'centered belief' in one game has predictive power for behaviour in another.

Before running the cluster analyses on participants' beliefs reported in the preceding Section, I ran three additional sessions of a team treatment aimed at incentivising reports of strategic reasoning, with a total of 76 participants. While the results are not very informative about the distinction between centered-belief players and other types, they provide an additional perspective on participants' reasoning. I therefore include a brief description on the data in the following final results Section.

5.6 Another glimpse into players' minds: team communication

Following Burchardi and Penczynski (2014), the game was played in pairs. In each pair, there is a suggesting player and a deciding player. The suggesting player suggests an action in the game and writes a free-form message that is

[16]A Boschloo test on whether choices below 17 are equally prevalent among centered-belief players and others yields $p = 0.03$. Note that I had to drop one out of 99 participants because due to a programming error, the participant managed to get past the screen of the 11–20 game without making a choice.

transferred to the deciding player alongside the suggested action. The deciding player sees both the suggestion and the message before deciding on the team's action, which gives the suggesting player an incentive to give a compelling argument for why the suggested action is a good choice.

Also following Burchardi and Penczynski (2014), I use the strategy method with respect to the player roles to save on participants (*i.e.*, both members of the team might be assigned either role and thus have to both supply a suggestion and a message and make a final decision after seeing their partner's message; importantly, along the path of action, communication remains one-way).[17] The actions in the game were 40, 44, and 48, which were to be split between the two team members in case the team was successful. At the end of the session, payoffs were converted into Euros at a rate of 2 points per Euro, and there was a show-up fee of 5 Euros.

As reported in Table 1, 29%, 26%, and 45% suggested 40, 44, and 48, respectively (the team decisions were 40, 44, and 48 with probabilities of 25%, 16%, and 59%). More importantly, though, 41% of those who suggest 40 explicitly mention "safety" in their messages. This consideration is virtually absent in others' explanations: only 4% of those who suggest 44 or 48 refer to "safety" (and another 4% to being "careful"). In contrast, 19% refer to "going for the risk" (although in more varied terms including "All in" or "No risk no fun"), which is absent in the former group's messages.

The notion of the lowest option being "safe" clearly suggests a belief that everybody else will be choosing either 44 or 48–which would correspond to a level-2 player in Level-k based on a level-0 that chooses by the option suggesting the highest payoff (Arad and Rubinstein's, 2012, argument for using "20" as the level-0 choice in their 11–20 game), or an NI-2 player in the noisy-introspection model. Note, however, that this observation does not square up with the numbers from the cluster analysis on reported beliefs: According to Table 3, clearly less than 20% belong to a belief cluster that would suggest "27" to be a "safe" choice—which, empirically, it is not.

6 Conclusion

In a certain type of games, this paper identifies a new type of players in the population that cannot be neglected, neither in terms of its prevalence nor in terms of its effect on the aggregate choice pattern—and, as a consequence, also on other players' empirical best-response in the game. I call the type of players 'centered-belief players' because they have a roughly-symmetric belief that peaks in the

[17] The treatment was preceded by a plain-vanilla public-good experiment, and only one of the two parts was randomly selected for payment.

centre of the option set. I call them 'heuristic' because their choices cannot be explained by either best- nor quantal-responding to their beliefs, in particular in the 11–20 game (which sets them apart from, *e.g.*, level-k or noisy-introspection types). Giving a 'centered belief' is not 'weird': it simply is a specific form of a central-tendency bias.[18] 'Centered beliefs' may come from a pre-conception that 'most variables lead to a Gauss curve' (think of age, IQ scores, …), which would likely be triggered only if the underlying variable has a natural ordering.

The 'heuristic centered-belief players' strongly affect the aggregate data pattern in the discoordination games and, therefore, also the expected payoffs from each action in the game. The 'heuristic' players' response to their belief may be driven by a reasoning of 'playing it safe' which, however, leads to making an empirically dominated choice. Seeing the 'heuristic centered-belief type' in two rather different games suggests that it is a general phenomenon to some degree. Seeing the predictive power of players' 'centered beliefs' for their actions in another game shows that a certain fraction of the population uses 'heuristic centered-belief play' as a general strategy within a certain type of games.

The findings are important because they a) help to understand the behaviour of a non-negligible share of the population, b) because we need to take the player type into account when making predictions for new settings, and c) because the findings show that we cannot look only at aggregate data when searching for the best model of behaviour. They show even that we cannot look only at individual choices: Arad and Rubinstein (2012) claim that the 11–20 game transparently and unambiguously identifies participants' levels of reasoning because "[i]t is hard to think of plausible alternative decision rules for this game. (...) The only other conceivable rules of behavior we could think of are randomly choosing a strategy or arbitrarily guessing the other player's strategy and best-responding to it." As I show in this paper, 'heuristic centered-belief players' neither choose their strategy "randomly" (at least not in the usual sense) nor do they best-respond to arbitrary guesses, and yet, they do follow an alternative decision rule. What alternative decision rule they follow, and exactly which features of a situation trigger 'heuristic centered-belief play' is an important topic of further research. As an immediate next step, the findings of Crosetto et al. (2020) suggest that first-price auctions might be another setting where 'heuristic centered-belief players' could play a role.

[18]'Centered beliefs' are a specific form of central-tendency bias because they require a somewhat symmetric belief in addition to the increased subjective-probability on central options; note that in the cluster analysis on beliefs in Section 5.3, there is an additional belief cluster that also peaks in the centre, but that is highly asymmetric.

Appendix

Robustness of the cluster analysis

Number of Clusters	1	2	3	4	5	6	7	8	9	10
Identified 'centered-belief types' (in %)	100	34	26	47	40	31	22	22	24	24

Table A.6: Proportion of participants classified as having 'centered beliefs' using cluster analyses for different numbers of clusters (using standard k-means clustering as inplemented in the kmeans function in R).

References

Apesteguia, J., Huck, S., and Oechssler, J. (2007). Imitation—theory and experimental evidence. *Journal of Economic Theory*, 136(1):217–235.

Arad, A. and Rubinstein, A. (2012). Multi-dimensional iterative reasoning in action: The case of the colonel blotto game. *Journal of Economic Behavior & Organization*, 84(2):571–585.

Bardsley, N., Mehta, J., Starmer, C., and Sugden, R. (2010). Explaining focal points: Cognitive hierarchy theory *Versus* team reasoning. *Economic Journal*, 120:40–79.

Bauer, D. and Wolff, I. (2018). Biases in beliefs: Experiment evidence. TWI Research Paper Series, No. 109.

Berger, U., De Silva, H., and Fellner-Röhling, G. (2016). Cognitive hierarchies in the minimizer game. *Journal of Economic Behavior & Organization*, 130:337–348.

Bock, O., Baetge, I., and Nicklisch, A. (2014). hroot: Hamburg registration and organization online tool. *European Economic Review*, 71:117–120.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2012). Salience theory of choice under risk. *The Quarterly Journal of Economics*, 127(3):1243–1285.

Burchardi, K. B. and Penczynski, S. P. (2014). Out of your mind: Estimating individual reasoning in one shot games. *Games and Economic Behavior*, 84:39–57.

Camerer, C. F., Ho, T.-H., and Chong, J.-K. (2004). A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 119(3):861–898.

Chmura, T., Goerg, S. J., and Selten, R. (2014). Generalized impulse balance: An experimental test for a class of 3×3 games. *Review of Behavioral Economics*, 1:27–53.

Cooper, D. J., Fatás, E., Morales, A. J., and Qi, S. (2018). Consistent depth of reasoning in level-k models. Working paper.

Crawford, V. P., Costa-Gomes, M. A., and Iriberri, N. (2013). Structural models of nonequilibrium strategic thinking: Theory, evidence, and applications. *Journal of Economic Literature*, 51(1):5–62.

Crosetto, P., Filippin, A., Katuščák, P., and Smith, J. (2020). Central tendency bias in belief elicitation. *Journal of Economic Psychology*, 78:102273.

Duch, M. L., Grossmann, M. R., and Lauer, T. (2020). z-tree unleashed: A novel client-integrating architecture for conducting z-tree experiments over the internet. *Journal of Behavioral and Experimental Finance*, 28:100400.

Eliaz, K. and Rubinstein, A. (2011). Edgar allan poe's riddle: Framing effects in repeated matching pennies games. *Games and Economic Behavior*, 71(1):88–99.

Faillo, M., Smerilli, A., and Sugden, R. (2017). Bounded best-response and collective-optimality reasoning in coordination games. *Journal of Economic Behavior & Organization*, 140:317–335.

Fischbacher, U. (2007). z-tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178.

Georganas, S., Healy, P. J., and Weber, R. A. (2015). On the persistence of strategic sophistication. *Journal of Economic Theory*, 159(A):369–400.

Goeree, J. K. and Holt, C. A. (2004). A model of noisy introspection. *Games and Economic Behavior*, 46(2):365–382.

Goeree, J. K., Louis, P., and Zhang, J. (2018). Noisy introspection in the 11–20 game. *The Economic Journal*, 128(611):1509–1530.

Haan, M. A., Lambert, S., and Winkel, B. M. (2009). Experimental results on collusion. pages 9–33. Cambridge University Press, Cambridge.

Hyndman, K., Terracol, A., and Vaksmann, J. (2022). Beliefs and (in)stability in normal-form games. *Experimental Economics*.

Mehta, J., Starmer, C., and Sugden, R. (1994). The nature of salience: An experimental investigation of pure coordination games. *American Economic Review*, 84(3):658.

Rubinstein, A. (2016). A typology of players: Between instinctive and contemplative. *The Quarterly Journal of Economics*, 131(2):859–890.

Selten, R. and Chmura, T. (2008). Stationary concepts for experimental 2×2-games. *The American Economic Review*, 98(3):938–966.

Sontuoso, A. and Bhatia, S. (2021). A notion of prominence for games with natural-language labels. *Quantitative Economics*, 12(1):283–312.

van Elten, J. and Penczynski, S. P. (2020). Coordination games with asymmetric payoffs: An experimental study with intra-group communication. *Journal of Economic Behavior & Organization*, 169:158–188.

Wolff, I. (2021). The lottery player's fallacy: Why labels predict strategic choices. *Journal of Economic Behavior & Organization*, 184:16–29.

THURGAU INSTITUTE
OF ECONOMICS
at the University of Konstanz

Hafenstrasse 6
CH-8280 Kreuzlingen

T +41 (0)71 677 05 10
F +41 (0)71 677 05 11

info@twi-kreuzlingen.ch
www.twi-kreuzlingen.ch