A diagnosis on the relationship between equilibrium play, stated beliefs, and best responses.*

Ciril Bosch-Rosa†
Thomas Meissner‡

December 10, 2017

Abstract

Experiments involving games have two dimensions of difficulty for subjects in the laboratory. One is understanding the rules and structure of the game and the other is forming beliefs about the behavior of other players. Typically, these two dimensions cannot be disentangled as belief formation crucially depends on the understanding of the game. We present a variation of the Two Player Guessing Game (Grosskopf and Nagel, 2008) which turns an otherwise strategic game into an individual decision-making task. This allows us to perform a within subject analysis of the decisions made for the same experiment, with and without strategic uncertainty. The results show that subjects with a better score at the individual decision making task form more accurate beliefs of other player's choices, and, additionally, better-respond to these beliefs. We also show that those who score higher at our individual task modify their beliefs based on the population they play against. This suggests that out of equilibrium play is mostly driven by a limited understanding of the game mechanics.

Keywords Guessing Game · Strategic Thinking · Cognitive Sophistication

JEL Classification C91 · D12 · D84 · G11

*Corresponding author: cirilbosch@gmail.com. The authors would like to thank Dan Friedman, Pedro Rey Biel, Luba Petersen, David Freeman, and especially Yves Breitmoiser for their comments. Additionally, the authors acknowledge financial support from the Deutsche Forschungsgemeinschaft (DFG) through CRC 649 "Economic Risk". The first author also would like to thank the Deutsche Forschungsgemeinschaft (DFG) for funding through the CRC TRR 190 "Rationality and Competition."

†Berlin University of Technology and Colegio Universitario de Estudios Financieros

‡Maastricht University

1 Introduction

Subjects in laboratory experiments consistently deviate from equilibrium behavior (Camerer, 2003). Many models of bounded rationality try to explain these deviations through errors in belief formation (e.g., Nagel (1995); Ho et al. (1998); Weizsäcker (2003)). A simpler explanation is that subjects may just not understand the rules of the game. Generally, when analyzing deviations from equilibrium behavior, one would expect both of these effects to play a role. However it is typically hard (if not impossible) to distinguish between the two, as correct belief formation crucially depends on a correct understanding of the rules of the game. With the help of a novel experiment, we are able to disentangle these two effects, thus improving the understanding of why subjects deviate from equilibrium behavior.

An extensive literature has attempted to analyze both belief formation and understanding the rules of the game. Costa-Gomes and Crawford (2006) present subjects with a series of two-player dominance-solvable games and conclude that most subjects understand the games, but play out of equilibrium solutions due to their "simplified models of others' decisions." In Costa-Gomes and Weizsäcker (2008) the authors look at subject's actions and their stated beliefs, and find that subjects rarely best respond to their stated beliefs. However, Rey-Biel (2009) observes that in simplified versions of the games studied in Costa-Gomes and Weizsäcker (2008), Nash Equilibrium is a better predictor of subject behavior than any other model based on level-K reasoning.

Another strand of the literature focuses on whether subjects understand the rules of the game. Using two-person guessing games, Chou et al. (2009) find that subjects are surprisingly unable to understand the experimental setup they are put in. By using different sets of instructions for the same game, and by introducing hints, they show that subjects do not deviate from equilibrium because of cognitive biases, but rather due to a lack of game form recognition, which they define as the relationships between possible choices, outcomes and payoffs.

In this experiment, we use a "one-player guessing game" which allows us to measure how well subjects understand the "mechanics" of the two-player guessing game (Grosskopf and Nagel, 2008). By comparing subjects' behavior in this "game" with behavior in a two-player guessing game, we can analyze to what extent the understanding of the rules of the game determines beliefs and their best-responses.[1]

The experimental results show that subjects who score higher at the one-player guessing game play values closer to the Nash Equilibrium in the two-player guessing game, and are better at best-responding to their beliefs. Additionally, we show that subjects who score high at the one-player guessing game change their beliefs according to the population they face. This suggests that out of equilibrium play is not only the result of a limited ability to form correct beliefs, but that it also results from the inability of subjects to fully understand and solve the game they are participating in.

2 Experimental Design

The experiment consists of four different parts: Subjects first play the one-Player guessing game (1PG), followed by the two-player guessing game (2PG). After this, we elicit their beliefs about other subject's two-player guessing game play. A subset of subjects then completed an additonal belief elicitaton task ("What-if" Belief Elicitation). At the end of the experiment, all subjects are asked to answer a battery of cognitive ability tests.

2.1 The One-Player Guessing Game (1PG)

The one-player guessing game was first introduced in Bosch-Rosa et al. (2017) and allows to test whether subjects can solve the two-player guessing game introduced by Grosskopf and Nagel (2008) *free of any strategic concerns.*

In essence, subjects play the role of both players in a two-player guessing game, i.e. they play the two player guessing game "against themselves". Accordingly, each subject (i) picks two numbers $x_i \in [0, 100]$ and $y_i \in [0, 100]$ and is paid depending on the absolute distance of each chosen number to the "target value" which is two thirds of the average of both numbers. The further away each chosen number is from this target value, the lower is the payoff. Formally the experimental payoff for choosing number x_i and y_i is:

$$\pi_i^{1PG}(x_i) = 1€ - 0.05€ \left| x_i - \frac{2}{3}\frac{x_i + y_i}{2} \right| \tag{1}$$

[1]For convenience we will henceforth refer to the one-player guessing game as a game, even though strictly speaking it is not one.

$$\pi_i^{1PG}(y_i) = 1€ - 0.05€ \left| y_i - \frac{2}{3}\frac{x_i + y_i}{2} \right| \tag{2}$$

Subjects are paid for both choices, so their combined payoff is:

$$\Pi_i^{1PG} = \max\left[\pi_i^{1PG}(x_i) + \pi_i^{1PG}(y_i), 0\right] \tag{3}$$

The payoff function is maximized at ($y_i = 0, x_i = 0$). This solution can be found through logical induction by starting with a random value $x_{0,i}$, and then calculating the "best response" which is $y'_{1,i} = \frac{1}{2}x_{0,i}$. Following this, a "best response to the best response" can be calculated ($x'_{1,i} = \frac{1}{2}y'_{1,i}$) and so on until reaching the fixed point ($x'_{\infty,i} = 0$, $y'_{\infty,i} = 0$).

By turning the two-player guessing game into an algebraic problem with no strategic uncertainty, we can separate those subjects who can solve the mathematical problem associated with the guessing game from those who cannot. Furthermore, our payoff structure allows us to rank every subject in the experiment according to their understanding of the game (i.e., their payoff in the 1PG).

2.2 The Two-Player Guessing Game (2PG)

The two-player guessing game that we use is an adaptation of the one presented in Grosskopf and Nagel (2008), and recently used in Nagel et al. (2016). Subjects are matched in pairs and asked to simultaneously pick a number $z_i \in [0, 100]$. In Grosskopf and Nagel (2008) the winner is whoever picks the number closer to 2/3 of the average of both numbers, so unlike the games with $N > 2$ subjects, now $z_i = 0$ is a (unique) weakly dominant strategy. In our version of the game, the payments are based on the (absolute) distance of each individual pick to 2/3 of the average of both numbers. Formally, the payment for player i depends on the choices of player j and her own in the following way:[2]

$$\Pi_i^{2PG} = \max\left[2€ - 0.10\left|\left(X_i - \frac{2}{3}\frac{z_i + z_j}{2}\right)\right|, 0\right] \tag{4}$$

This small change in payoffs dramatically changes the game as now the equilibrium is reached through iterated deletion of *weakly dominated* strategies, and zero is no longer a

[2]Note that we limited the minimum payoff to zero in order to avoid potential losses for the subject.

weakly dominant strategy. Now the best response is to choose 1/2 of the number a player thinks its counterpart chooses.

We opted for this modification of the original game for two reasons. First, it allows us to *de facto* ask subjects for a point estimate of their belief about the other subject's choice, and secondly, and more important, it makes the game comparable to the 1PG.

2.3 Belief Elicitation

After subjects had played the 1PG and the 2PG (with no feedback in both cases) we elicited their beliefs about the other players'decisions in the two-player guessing game. Similar to Lahav (2015), subjects were asked to distribute a total of 19 "tokens" into 20 "bins".

Each token represented a subject in the session (each session consisted of 20 subjects), and each bin had a range of 4 integers that players could play in the 2PG (i.e. the first bin had the range [0,4], the second [5,9], and so on).[3].

To incentivize subjects, we used a linear scoring rule that paid €0.10 for each token that overlapped with the choice of any other subject in the 2PG.[4] Formally, define Π_i^B as the payoff for subject i, b_{ij} as the number of tokens that she deposited in bin j, and p_{-ij} as the number of subjects other than player i that chose a value that falls within bin j in the 2PG. Then the payoff for belief formation for subject i is:

$$\Pi_i^B = \sum_{j=1}^{20} 0.10€ * \min\left[b_{ij}, p_{-ij}\right] \tag{5}$$

.

The resulting distribution of beliefs provides an estimate of what subjects think about other subjects' choices, and also allows to analyze how capable subjects are to best respond to their own beliefs.[5]

[3]Figure 10 in AppendixC contains a screenshot of the experiment interface used to elicit beliefs

[4]For instance, if a subject put three tokens in the bin "5-9" and in her session only 2 subjects had actually played *any* value within this range, then she would get paid a total of 20 cents for the tokens allocated in that bin. If, on the other hand, she placed 5 token in the bin "0-4" and 10 subjects had played a value in this range, then she would be paid 50 cents for the tokens allocated in that bin.

[5]There is some discussion about how to best incentivize subjects to state their true beliefs. In particular, there is mixed evidence on whether incentive compatibility matters or not (Schotter and Trevino, 2014). Methods to elicit beliefs beyond first moments, such as ours, are typically difficult for subjects to

2.3.1 "What if" Belief Elicitation

Since there could be some influence of having played the 1PG on the beliefs about the 2PG plays, we additionally elicited beliefs about 2PG plays from subjects who had not played the 1PG before. Again, subjects were asked to distribute 19 tokens across the strategy space. The only difference was that, this time, subjects were asked to guess the choices of 19 subjects that had played the 2PG "a couple of weeks ago, *without having previously played the 1PG*". Beliefs were incentivized as above. A subset of 40 subjects participated in this task.

2.4 Cognitive Ability

Gill and Prowse (2016) show that subjects who score higher in a Raven Test (Raven, 1960) choose numbers closer to equilibrium, earn more, and converge quicker to equilibrium in a three-player guessing game. Since we are interested in studying the ability of subjects to solve the guessing game, we also tested the cognitive ability of our subjects. In particular, all subjects answered a Raven Test and played "Race-to-60," a variant of the Race game (see e.g. Gneezy et al. (2010), Levitt et al. (2011)).[6] The Raven Test is a multiple choice test in which subjects must pick an element that best completes a missing element in a matrix of geometrical shapes.[7] The score of this test has been found to correlate with measures of strategic sophistication and the ability of subjects to solve novel problems (Carpenter et al. (1990)). It is increasingly used in economic research due to its simplicity and the lack of required technical skills.[8]

Since logical induction is a central element of the guessing games, we test this ability with the "Race-to-60" game. In this game, each participant and a computerized player sequentially choose numbers between 1 and 10, which are added up. Whoever is first

understand. Using a non-linear scoring rule would introduce an additional level of complexity. Hence, while not incentive compatible for risk neutral subjects, we opted for this approach because we believe it provides the best compromise between tractability for subjects and incentivization efficacy.

[6] While in Gill and Prowse (2016) subjects go through all 60 matrices of the original Raven Test, in our case subjects just took part in three of the hardest blocks of 12 matrices.

[7] Figure 11in AppendixC contains a screenshot of the experiment interface.

[8] Another important reason why we prefer the Raven test over alternatives such as the Cognitive Reflection Test (Frederick (2005)) is that we have found out that many of our subjects have seen the CRT outside of the lab, while the Raven test is not as known. Additionally, we hypothesize that previous experience with Raven matrices distorts the measure less than previous experience with the CRT.

to push the sum to or above 60 wins the game. The game is solvable by backward induction, and the first mover can always win.[9] Subjects always move first and therefore, independently of the computer strategic sophistication, they can always win the game.[10]

3 Results

A total of 80 subjects participated in this experiment. All subjects were recruited through ORSEE (Greiner, 2004) and were mostly undergraduate students with a variety of backgrounds, ranging from anthropology to electrical engineering or architecture. Sessions lasted one and a half hours and were run at the Experimental Economics Laboratory of the Berlin University of Technology. The experiment was programmed and conducted with the software z-Tree (Fischbacher, 2007). All results are listed in this section except the cognitive ability results which can be found in Appendix A.

3.1 The One Player Guessing Game

In Figure 1 we present the results of the 1PG in a scatter plot. Recall that in this case subjects have to pick two numbers, (x_i, y_i); the first number is depicted on the horizontal axis, the second on the vertical axis. In light blue we show those subjects that fully solved the game (0,0). The diagonal dashed line marks the points where a subject picked the same number for x_i and y_i. This is an important indicator as in this task there are two ways in which a subject (that has not fully solved the game) can improve her payoffs; by picking numbers closer to zero, and/or by picking numbers that are closer to each other.[11]

In Figure 1 we cans see how subjects who play numbers closer together, also play numbers closer to the origin. A Spearman test confirms the correlation between higher average of both choices and the distance between them (Spearman ρ=0.83, p-value <0.001). We can therefore confidently use the payoffs of the 1PG as a measure of the understanding

[9] By picking numbers such that the common pool adds up to the sequence : [5; 16; 27; 38; 49; 60] the first mover can always win this game.

[10] As in Bosch-Rosa et al. (2017) the backward induction ability of the computer increased with the rounds of the game. In the first round the computer could do only one step of backward induction, in the second it could do two, in the third three, and so on.

[11] An example for the first case would be picking twice 66, or twice 33. In the former case the subject would get €0, in the second €.90. For the second case, imagine a subjects has to decide between the pair (10, 56), and the pair (32, 34), both average 33, but the former pays €0.35 and the latter €0.90.

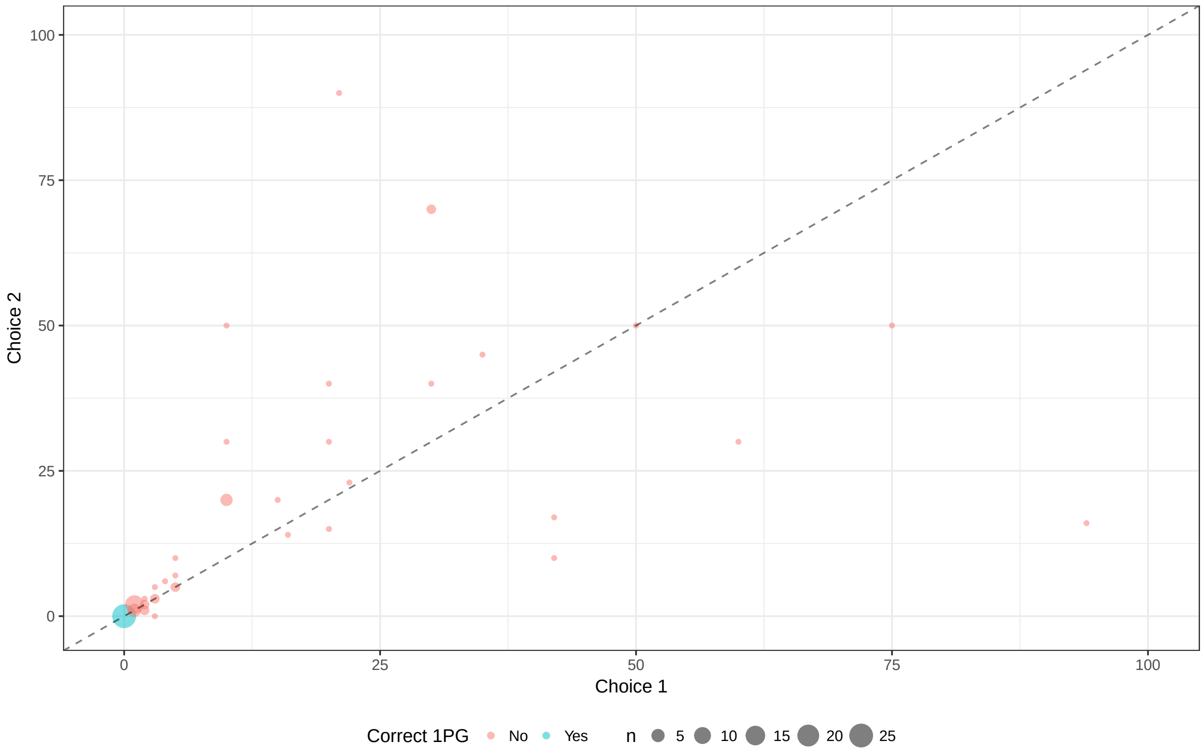

Figure 1: Scatter plot of the choices made by each subject. In light blue the choices made by those that fully solved the game (0,0). See Figure 13 in Appendix C for a zoom in plot depcting only the choices from [0,50].

of the structure of the guessing game, as those subjects with high payoffs played both numbers that were close to each other, and to zero. Therefore, from now on, we will use the payoff in the 1PG as a measure of "understanding" of the mechanics of the guessing game.

Surprisingly, only a minority (≈31%) of subjects is able to fully solve the task, i.e. pick zero for both numbers.

Result 1: *Only 31% of our subjects fully understand the one-player-guessing game.* While not directly comparable, these numbers are better than those of Grosskopf and Nagel (2008) where only 10% of their subjects play the (weakly) dominant strategy of choosing zero. The authors hypothesized that subjects performed poorly because they were either not realizing how much influence they had in the game, or were trying to find a fixed point. In our case looking for a fixed point is precisely what subjects should do, and it is unequivocal that they are in full control of the game. Therefore, neither of their reasons can explain our subjects' poor performance. Note, however, that if we relax the criteria of what it means to "understand" the task to picking both numbers less

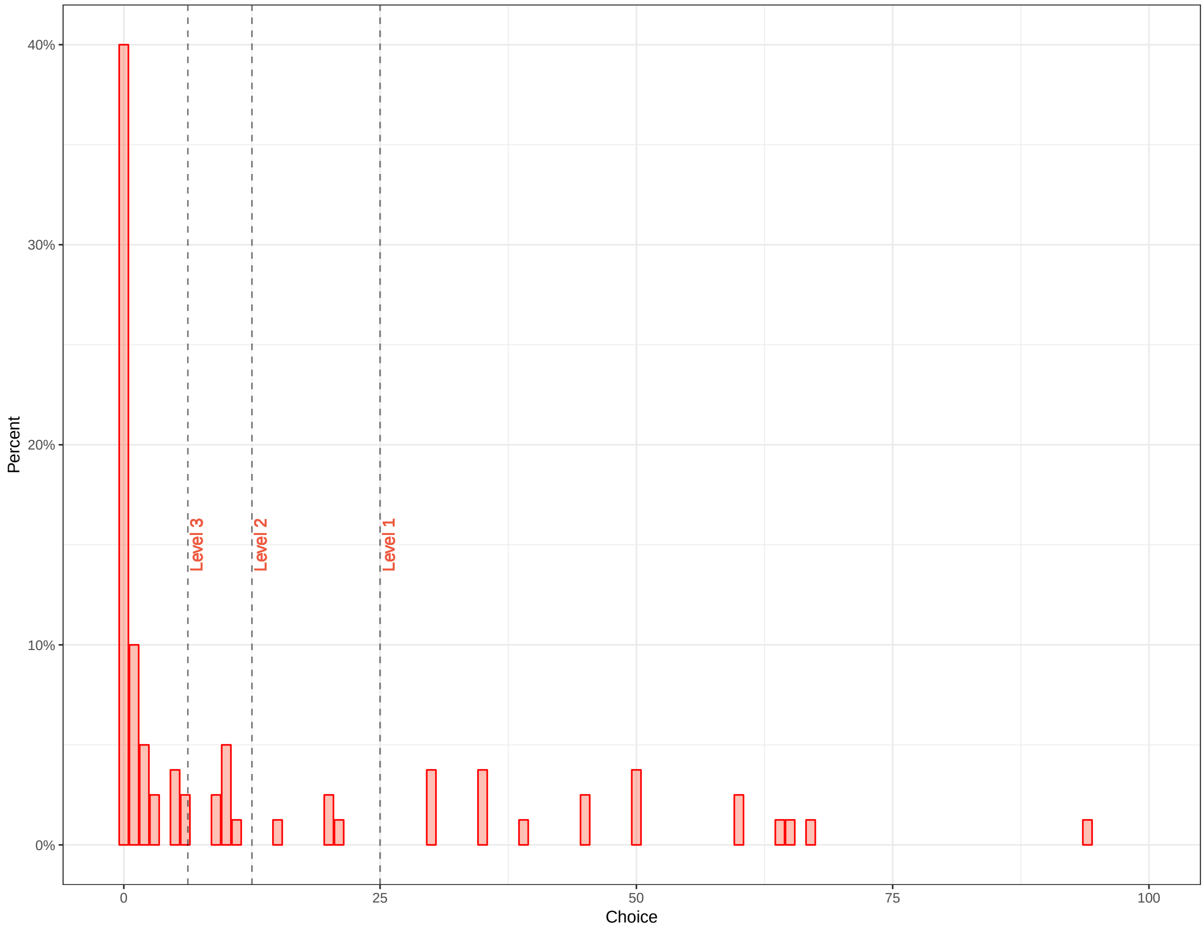

Figure 2: Distribution of choices in the 2PG.

than two, then 30 subjects do so which increases to 37.5% the percentage of subjects that "understood" the game.

3.2 The Two Player Guessing Game

Figure 2 shows the distribution of choices played in the 2PG. The distribution appears to be rather different compared to the typical distribution one sees with guessing game "first timers." The mass of the distribution is close to zero with 40% of subjects playing the Nash Equilibrium. The mean is of 13.47 and the median choice is 2. These low numbers could be the result of two phenomena: introspective learning from having played the 1PG, a change in beliefs normally formed when playing against subjects with no experience given that now the pool of subjects has experience. In section 3.4.1 we show that these low numbers are, in general, a consequence of introspective learning and not a shift in beliefs. So while most subjects are not able to fully solve the 1PG, there are

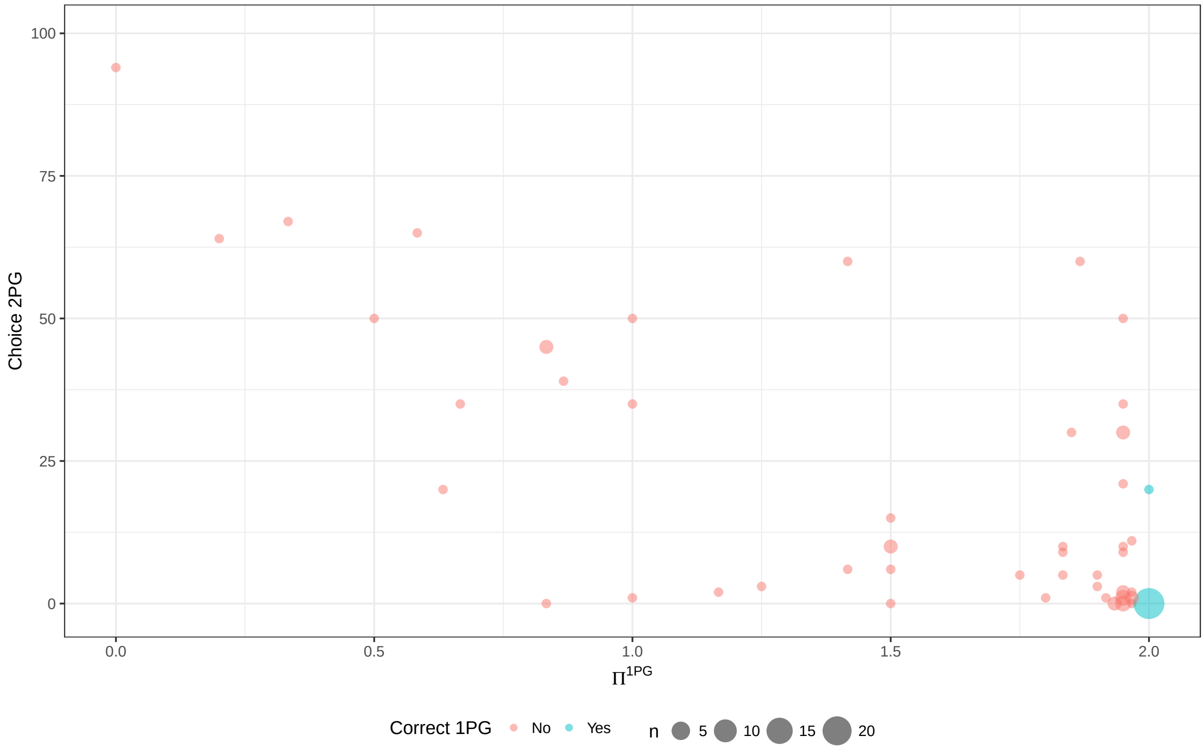

Figure 3: Distribution of choices in the 2PG (vertical axis) and payoff in the 1PG (horizontal axis).

large learning effects derived from thinking about it. Interestingly, the distribution has no typically observed "spikes" at the values 25 and 12.5. These values have significance as they would result from level-1 and level-2 reasoning respectively.[12]

3.3 Relationship between the 1PG and the 2PG

Figure 3 shows the decisions of subjects in the 2PG on the vertical axis, and their payoffs for the 1PG on the horizontal axis. The negative correlation between both is apparent: subjects who earn higher payoffs in the 1PG play lower numbers in the 2PG (Spearman ρ=0.74, p-value <0.001). Moreover, of the subjects who fully solved the 1PG (light blue dots in Figure 3), 96% (24/25) chose zero in the 2PG.

Result 2: *Subjects with a better understanding of the structure of the game (i.e. higher payoffs in the 1PG) play numbers closer to the Nash Equilibrium.*

But, is playing numbers near the Nash Equilibrium the best strategy in the 2PG? To answer this question we construct $\bar{\Pi}_i^{2PG}$. This variable represents the payoff that each

[12] Notice that these values are different from the usual 33 and 22 in usual Guessing Games, as in the Two Player Guessing Game subjects need to take into account the weight that their own choice carries in the final outcome.

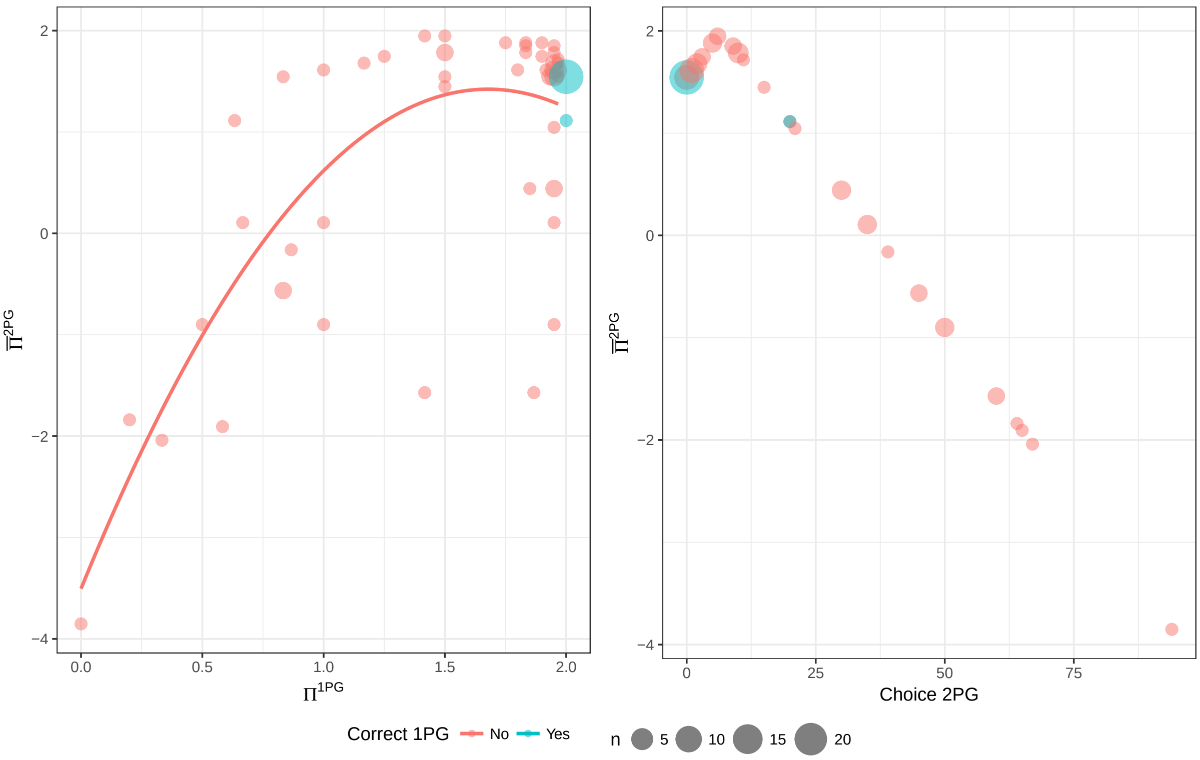

Figure 4: Relationship of payoff in the 1PG (Π_i^{1PG}) and $\bar{\Pi}_i^{2PG}$ (left panel). The red line is a fitted quadratic function. The right panel shows the relationship of choice in the 2PG (z_i) and $\bar{\Pi}_i^{2PG}$. In both panels the light blue dots refer to subjects that fully solved the 1PG.

subject i would have gotten had she played against the average choice of all other subjects j except herself (i.e., $j \neq i$). Formally $\bar{\Pi}_i^{2PG}$ is defined as:

$$\bar{\Pi}_i^{2PG} = 2€ - 0.10 \left| z_i - \frac{2}{3} \frac{z_i + \sum_{j \neq i}^{N} z_j/(N-1)}{2} \right| \tag{6}$$

Figure 4 illustrates the relationship of $\bar{\Pi}_i^{2PG}$ with both, the payoffs of the 1PG ($\Pi^1 PG_i$, left panel) and choice in the 2PG (z_i, right panel). Both graphs show a non-monotonic pattern (see Figure 14 in Appendix C for a close-up of Figure 4). On the left we see how subjects who performed very poorly in the 1PG have a low $\bar{\Pi}_i^{2PG}$. At the the same time, subjects who fully solved the 1PG don't have the highest $\bar{\Pi}_i^{2PG}$. This is because they play the Nash Equilibrium, when payoffs would have been maximized by playing a number close to 9 as can be seen in the right plot. Regressing $\bar{\Pi}_i^{2PG}$ on $\Pi^1 PG_i$ and $(\Pi^1 PG_i)^2$ we get a coefficient which is significantly different from zero and negative for the squared term. This gives statistical support to the fitted quadratic function (red line) in Figure 4.

Result 3: *The relationship between ability to solve the 1PG and payoffs in the 2PG follows a non-monotonic pattern.*

	$\bar{\Pi}_i^{2PG}$
Π_i^{1PG}	5.519***
	(0.904)
$(\Pi_i^{1PG})^2$	-1.565***
	(0.353)
Constant	-3.389***
	(0.500)
N	80
adj. R^2	0.580
Joint test p-value	0.000

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 1: Regression of $PayOS_i$ on $PaymentSelf_i$ and $PaymentSelf_i^2$.

3.4 Subjective Beliefs

To analyze whether ability to understand the structure of the game (measured by the payoff in the 1PG) influences belief formation, we plot Π_i^{1PG} against the number of correctly guessed tokens in the left panel of Figure 5. Again, in order to avoid noise due to session specific outliers, we compute the number of correct tokens by comparing individual beliefs to the distribution of 2PG choices over all sessions (for more details see Appendix B). A clear pattern arises in which subjects who score higher in the 1PG get guess a larger number of correct tokens (Spearman ρ=0.583, p-value<0.000).

Result 4: *Subjects with a better understanding of the structure of the game (i.e., higher payoffs in the 1PG) have more accurate beliefs.*

On the right panel of Figure 5 we plot the distribution of tokens (horizontal axis) against the payoff in the 1PG (vertical axis). While subjects who score low in the 1PG spread out their tokens across most of the strategy space, subjects who performed better expect their counterparts to play numbers closer to the Nash Equilibrium (Spearman ρ=-0.513, p-value<0.001).

An interesting question is whether choices in the 2PG are best responses to the stated beliefs (i.e., the token distribution). To analyze this question we compute the choice in the 2PG that would maximize the payoff of a subject conditional on her stated beliefs being correct (Equation 7):

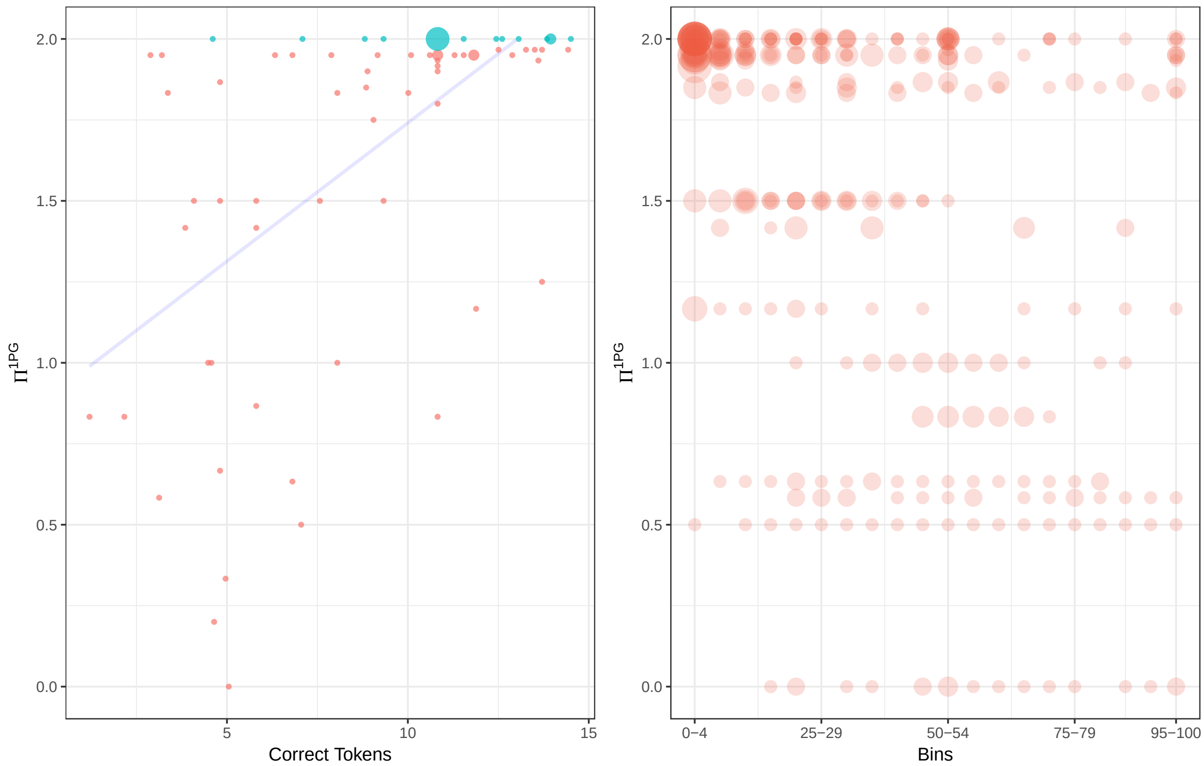

Figure 5: Payoff in the 1PG vs. number of correct tokens (left panel) and number of tokens assigned to the different bins (right panel). The purple line on the left panel is a linear fit to the data.

$$z_i^*(B) = \arg\max_{\tilde{z}_i} \sum_{j=1}^{20} \frac{b_{ij}}{19} \left[2 - 0.10 \left| \tilde{z}_i - \frac{2}{3} \frac{C + \bar{b}_j}{2} \right| \right], \tag{7}$$

where $z_i^*(B_i)$ is the choice of subject i that maximizes her payoffs given her beliefs $B_i = (b_{i1}, b_{i2}, ..., b_{i20})$, b_{ij} is the number of tokens that subject i put in bin j, and $\bar{b}_j$ is the average value of the bin (so for example, for the first bin [0,4], $\bar{b}_1 = 2$, for the second [5,9], $\bar{b}_2 = 7$, etc.).[13] We then create an individual variable $\Delta z_i^* = |z_i - z_i^*(B_i)|$ which is the absolute difference between actual choice of subject i in the 2PG minus the optimal choice conditional on her stated beliefs. Figure 6 illustrates the relation of Δz_i^* and the payoffs for the 1PG (left panel) and 2PG (right panel). It is clear that the lower Δz_i^*, the higher Π_i^{1PG} (Spearman ρ=-0.27, p-value=0.015). Similarly, the lower is Δz_i^* the higher is $\bar{\Pi}_i^{2PG}$ (Spearman ρ=-0.32, p-value=0.003). These results imply that those subjects that scored higher in the 1PG (i.e., that have a better understanding of the guessing game mechanics) are better able to best respond to their own beliefs. Additionally we find that

[13]We pick this instead of the lowest value of the bin, because it is a more stringent test to our "high ability" subjects.

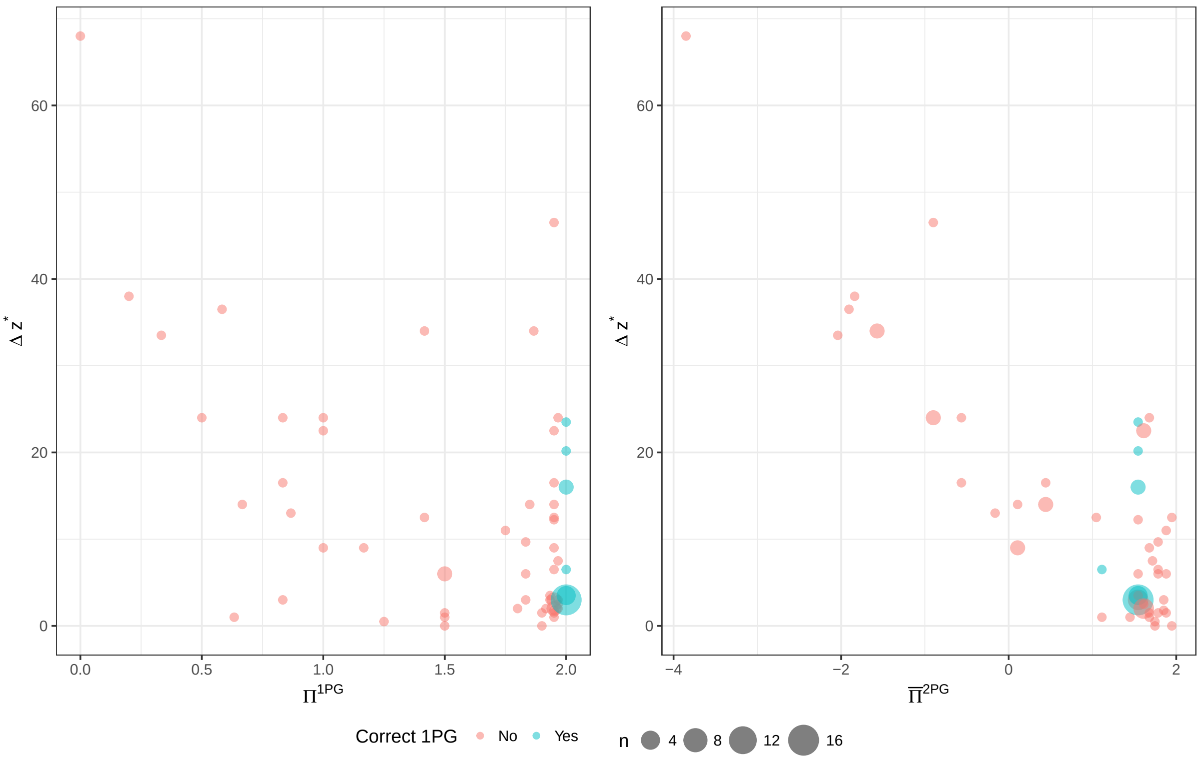

Figure 6: Difference between actual choice and optimal choice conditional on beliefs (Δz_i^*) vs. payoff in the 1PG (Π_i^{1PG}, left panel) and payoff against the average choice in the 2PG ($\bar{\Pi}^{2PG}$, right panel). In both panels the light blue dots refer to subjects who fully solved the 1PG.

more accurate best responses to own beliefs result in higher payoffs in the strategic 2PG.

Result 5: *Subjects with a better understanding of the structure of the game (i.e., higher payoffs in the 1PG) choose numbers closer to the best response of their beliefs and get a higher payoff in the 2PG.*

3.4.1 "What-if" Beliefs

Because there could be some influence of having played the 1PG on the beliefs for the 2PG we asked 40 subjects to guess the choices of 19 subjects that had played the 2PG "a couple of weeks ago, *without having previously played the 1PG*". From now on we will refer to this distribution as the WI distribution.

Figure 7 illustrates the token allocations for both the original belief elicitation task and the WI task. The differences are small, with the WI distribution slightly shifted away from the Nash Equilibrium. Yet, a different picture arises when looking at the data at an individual level. Comparing the mean of the first belief elicitation task $\bar{B}_{B,i}$) to the mean of the WI distribution ($\bar{B}_{W,i}$) a Wilcoxon matched-pairs signed-ranks test

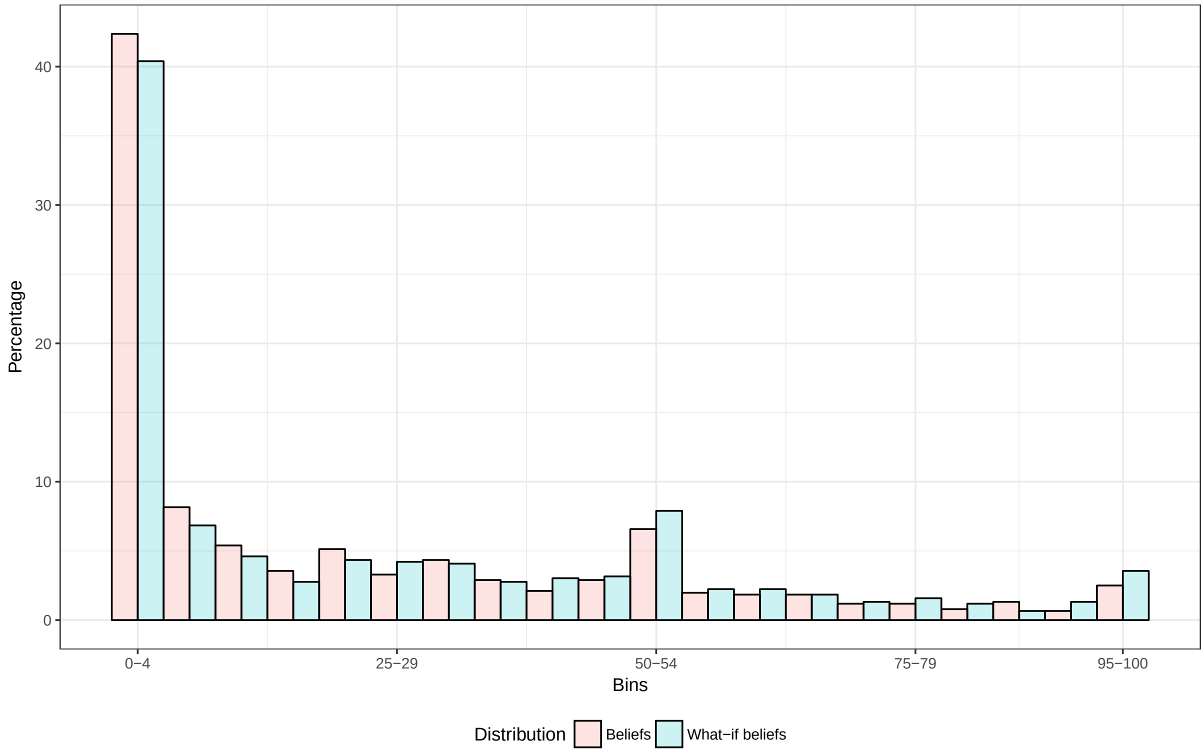

Figure 7: Distribution of beliefs and "what-if' beliefs.

shows differences within subject (p-value=0.039), this same test also shows individual level differences across the variances of both distributions (p-value¡0.01). Furthermore, if we use a one-sided Sign test of matched pairs for the difference in means and variance ($\Delta B_i = \bar{B}_{W,i} - \bar{B}_{B,i}$) and $\Delta\sigma_i^2 = \sigma_{W,i}^2 - \sigma_{B,i}^2$ respectively), we see how for both cases the values are positive and different from zero (p-value/¡0.02). In Appendix C, Figure 15 shows how most subjects shift the mean of their distribution away from the Nash Equilibrium, while increasing the variance in the WI distribution.

In Figure 8 we plot the change in mean (ΔB_i) in the vertical axis, and the payoff for the 1PG (Π^{1PG}) in the horizontal axis. Any value above the horizontal dotted line indicates a relative shift away from the NE for the WI distribution. As it is clear from the graph, the correlation between high payoff in the 1PG and a positive ΔB_i is positive and close enough to significance (p-value=0.0503). On the other hand this correlation seems not to be fully monotonic as the Spearman rank test is not significant (Spearman ρ=0.262, p-value=0.101).

Result 6: *There is a weak correlation between high payoffs in the 1PG and shifting beliefs for the "what-if" case away from the Nash Equilibrium.*

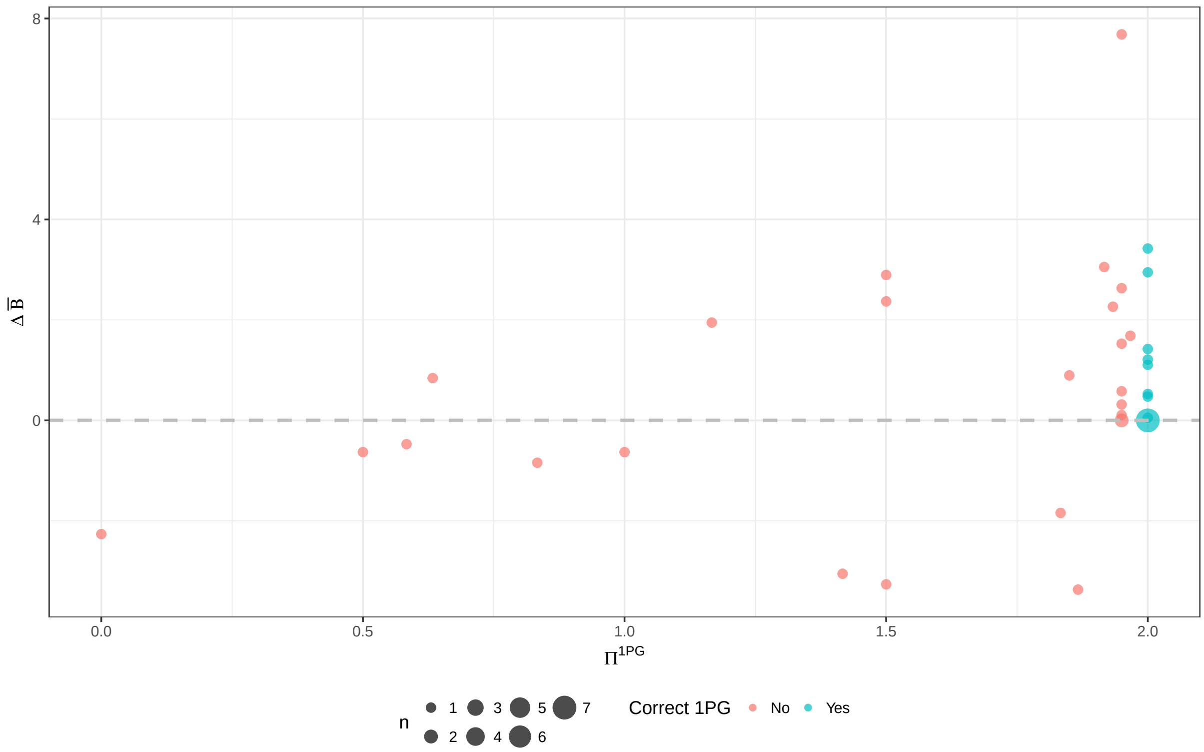

Figure 8: Payoff in the 1PG in the horizontal axis, and change in the mean between belief distributions (Δb_i) in the vertical axis. Any value of Δb_i) above zero is a change in the mean away from the Nash Equilibrium.

4 Conclusion

In laboratory experiments, subjects usually deviate from Nash Equilibrium. These deviations can be the result of either subjects not understanding the game setup correctly or from not forming the correct beliefs about the strategies of their counterparts. One strand of the literature has tried to explain these deviations as errors in belief formation (e.g., Costa-Gomes and Crawford (2006); Ho et al. (1998)). Yet, some recent research shows that subjects might not be fully understanding the experimental environment. Weizsäcker (2003) suggests that using less abstract environments might help subjects build better beliefs, while Chou et al. (2009) show that subjects appear to have little understanding of the two-player guessing game.

In this paper we use an individual decision-making task that allows uncoupling subject understanding of the mechanics of the game from her belief formation. Disentangling these two dimensions, we can establish the extent that each one of them contributes to the out-of-equilibrium play in experimental setups.

We find that subjects who have a better understanding of the game play closer to the Nash Equilibrium, are better at best-responding to their own beliefs, and modify their beliefs (correctly) depending on the pool of subjects they are facing. This result is inconsistent with models of the Level-k type (e.g., Costa-Gomes and Crawford (2006)) which assume that agents fully understand the mechanics of the game and only play out of equilibrium due to flaws in their belief formation. Our findings suggest, otherwise, that out of equilibrium play is not only the result of a limited ability to form correct beliefs, but that it also results from the inability of subjects to fully understand and solve the game they are participating in.

References

Bosch-Rosa, C., T. Meissner, and A. Bosch-Domènech (2017): "Cognitive bubbles," *Experimental Economics.* Cited on pages 3 and 7.

Camerer, C. (2003): *Behavioral game theory: Experiments in strategic interaction*, Princeton University Press. Cited on page 2.

Carpenter, P. A., M. A. Just, and P. Shell (1990): "What one intelligence test measures: a theoretical account of the processing in the Raven Progressive Matrices Test." *Psychological review*, 97, 404. Cited on page 6.

Chou, E., M. McConnell, R. Nagel, and C. R. Plott (2009): "The control of game form recognition in experiments: Understanding dominant strategy failures in a simple two person "guessing" game," *Experimental Economics*, 12, 159–179. Cited on pages 2 and 16.

Costa-Gomes, M. A. and V. P. Crawford (2006): "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review*, 96, 1737–1768. Cited on pages 2, 16, and 17.

Costa-Gomes, M. A. and G. Weizsäcker (2008): "Stated Beliefs and Play in Normal-Form Games," *The Review of Economic Studies*, 75, 729–762. Cited on page 2.

Fischbacher, U. (2007): "z-Tree: Zurich toolbox for ready-made economic experiments," *Experimental economics*, 10, 171–178. Cited on page 7.

Frederick, S. (2005): "Cognitive reflection and decision making," *The Journal of Economic Perspectives*, 19, 25–42. Cited on page 6.

Gill, D. and V. Prowse (2016): "Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis," *Journal of Political Economy*, 124, 1619–1676. Cited on page 6.

Gneezy, U., A. Rustichini, and A. Vostroknutov (2010): "Experience and insight in the race game," *Journal of Economic Behavior & Organization*, 75, 144–155. Cited on page 6.

GREINER, B. (2004): "The online recruitment system orsee 2.0-a guide for the organization of experiments in economics," *University of Cologne, Working paper series in economics*, 10, 63–104. Cited on page 7.

GROSSKOPF, B. AND R. NAGEL (2008): "The two-person beauty contest," *Games and Economic Behavior*, 62, 93–99. Cited on pages 1, 2, 3, 4, and 8.

HO, T.-H., C. CAMERER, AND K. WEIGELT (1998): "Iterated dominance and iterated best response in experimental" p-beauty contests"," *The American Economic Review*, 88, 947–969. Cited on pages 2 and 16.

LAHAV, Y. (2015): "Eliciting beliefs in beauty contest experiments," *Economics Letters*, 137, 45–49. Cited on page 5.

LEVITT, S. D., J. A. LIST, AND E. SALLY (2011): "Checkmate: Exploring backward induction among chess players," *The American Economic Review*, 101, 975–990. Cited on page 6.

NAGEL, R. (1995): "Unraveling in guessing games: An experimental study," *The American Economic Review*, 85, 1313–1326. Cited on page 2.

NAGEL, R., C. BÜHREN, AND B. FRANK (2016): "Inspired and inspiring: Hervé Moulin and the discovery of the beauty contest game," *Mathematical Social Sciences*. Cited on page 4.

RAVEN, J. C. (1960): *Guide to the standard progressive matrices: sets A, B, C, D and E*, HK Lewis. Cited on page 6.

REY-BIEL, P. (2009): "Equilibrium play and best response to (stated) beliefs in normal form games," *Games and Economic Behavior*, 65, 572–585. Cited on page 2.

SCHOTTER, A. AND I. TREVINO (2014): "Belief elicitation in the laboratory," *Annu. Rev. Econ.*, 6, 103–128. Cited on page 5.

WEIZSÄCKER, G. (2003): "Ignoring the rationality of others: evidence from experimental normal-form games," *Games and Economic Behavior*, 44, 145–171. Cited on pages 2 and 16.

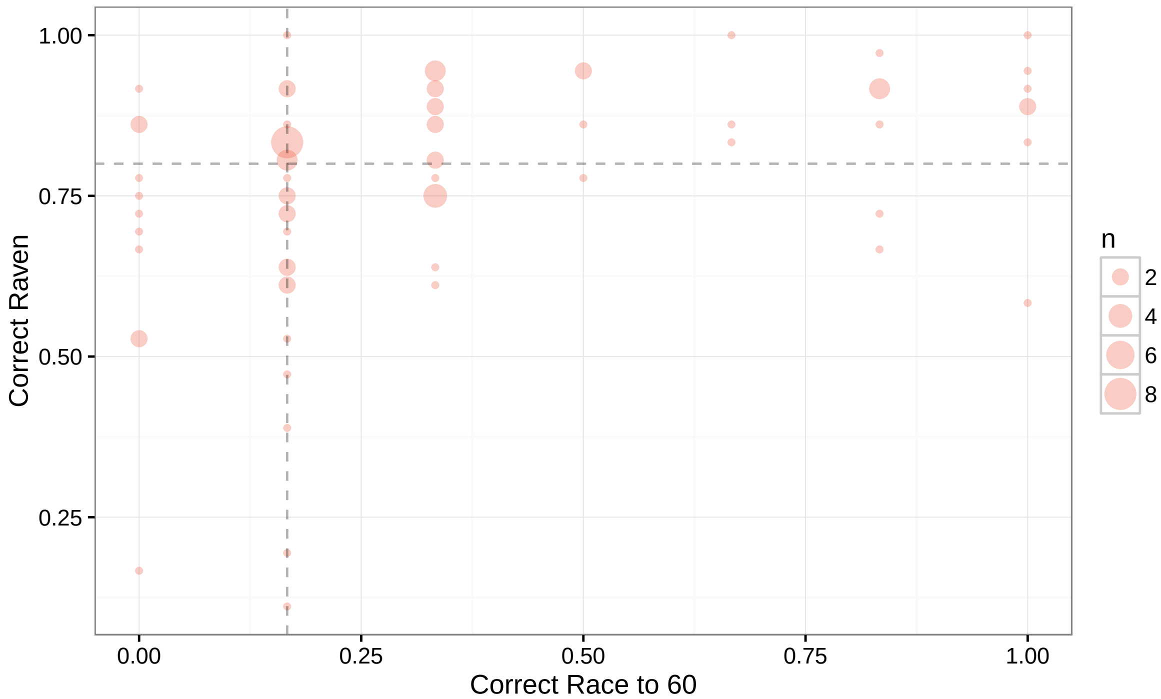

Figure 9: Normalized results for Race to 60 (horizontal axis) and Raven Test (vetical axis). The dashed lines mark the median result for each test.

A Cognitive Ability

As described in section A all subjects were asked to solve thirty-six matrices from the Raven Test, and took part in six rounds of Race-to-60. The normalized results of these two measures are plotted in Figure 9. Most subjects performed well at the Raven test (median of twenty nine solved matrices), but not in the Race-to-60 game (median of one won round). Interestingly, the lower right quadrant of the scatter plot is completely empty. This implies that while a good score in the Race-to-60 game requires a good score in the Raven Test, the opposite is not true. The correlation between both measures of cognitive ability is positive and significant, as well as with the payments for the 1PG and 2PG (see Table 2).

	Raven	Race	Payment 1PG	Payment 2PG	Difference	Change
Raven	1					
Raven	0.290***	1				
Payment 1PG	0.225***	0.321***	1			
Payment 2PG	0.235***	0.294***	0.344***	1		
Difference	-0.228***	-0.222***	-0.407***	-0.523***	1	
Change	-0.051	0.116	0.000	0.334***	-0.136	1

Table 2: Correlations

	(1) Correct 1PG	(2) Π_i^{1PG}	(3) $\bar{\Pi}_i^{2PG}$	(4) Δz_i^*
Raven	0.706**	0.390	1.217	-9.371
	(0.279)	(0.326)	(0.756)	(8.084)
Race	0.505***	0.548***	0.814*	-9.594**
	(0.163)	(0.190)	(0.440)	(4.708)
Constant	-0.409*	1.176***	-0.172	20.90***
	(0.208)	(0.243)	(0.563)	(6.014)
N	80	80	80	80
adj. R^2	0.226	0.133	0.086	0.074

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Linear regression of cognitive ability measures on Correct 1PG, Π_i^{1PG}, $\bar{\Pi}_i^{2PG}$, Δz_i^* and ΔB_i.

In the first column of Table 3 we report a linear probability model where the binary dependent variable is having solved the 1PG. In the second column, we report a linear regression with the payment in the 1PG as the dependent variable.[14] In both cases, the independent variables are the (normalized) cognitive ability test scores. The results show an interesting asymmetry; subjects that fully solve the 1PG scored high both at the the Race-to-60 and the Raven test, while only backward induction seems to have an effect on the payoff for the 1PG.[15]

Result: *Both a high score at the Raven test and Race-to-60 game are important to fully solve the 1PG, while only Race-to-60 seems to matter for the 1PG payoff.*

Additionally, in column 3 $\bar{\Pi}_i^{2PG}$ is regressed on both measures of cognitive ability, and only Race-to-60 is marginally important to score high in the 2PG. In column 4 Δz_i^* is the dependent variable, and this time the coefficient for Race-to-60 is highly significant and negative. This means that backward induction abilities are central for a subject to be able to best respond to her own beliefs. On the other hand, in column 5 we see how neither performing well at the Raven test or at the Race-to-60 game imply a higher measure of ΔB_i.

[14] A Tobit model for the 1PG payoff gives the same results.

[15] The Variance Inflation Factor discards any severe case of multicollinearity between both regressors.

B Normalization

The way in which we create the "generic" distribution for each subject is by adding all 2PG choices (except hers) for each bin that were played in her order. This results in a distribution of 79 choices across the 20 possible bins. We then normalize this to look like a distribution made by 19 subjects by multiplying the resulting number in each bin by 19/79.

C Extra Graphs and Figures

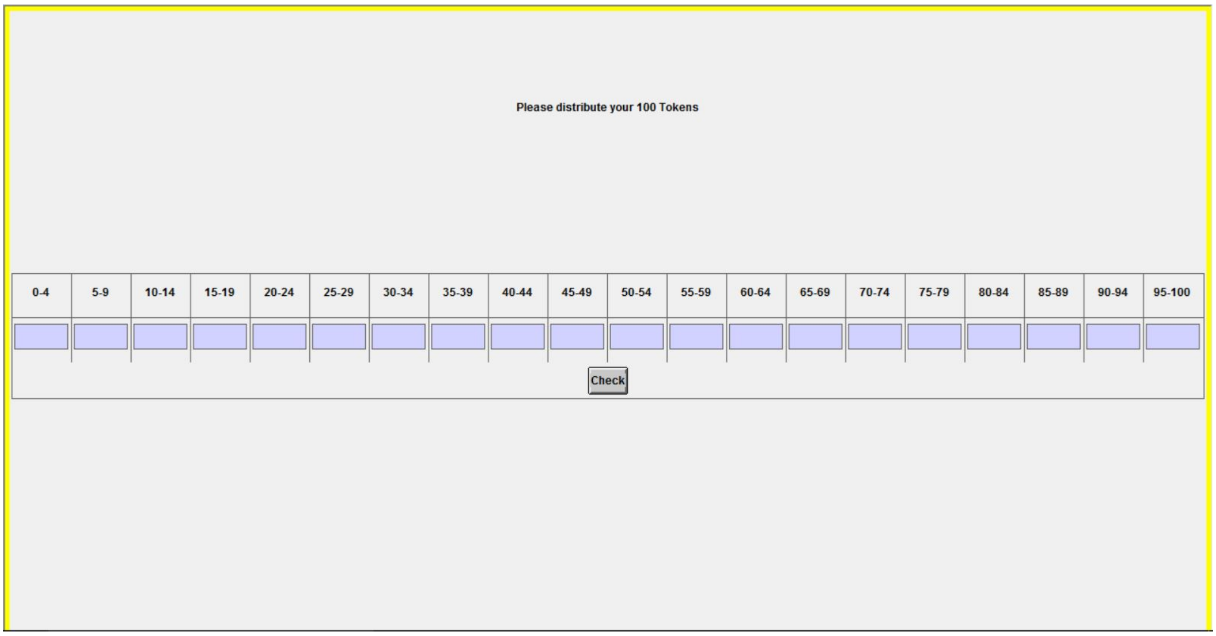

Figure 10: Distribution Screen. Subjects were provided with N-1 tokens to distribute across the 20 bins provided in the screen. The check button counted for them the amount of tokens they had deposited at any moment. Subjects could use less than 19 tokens, but never more.

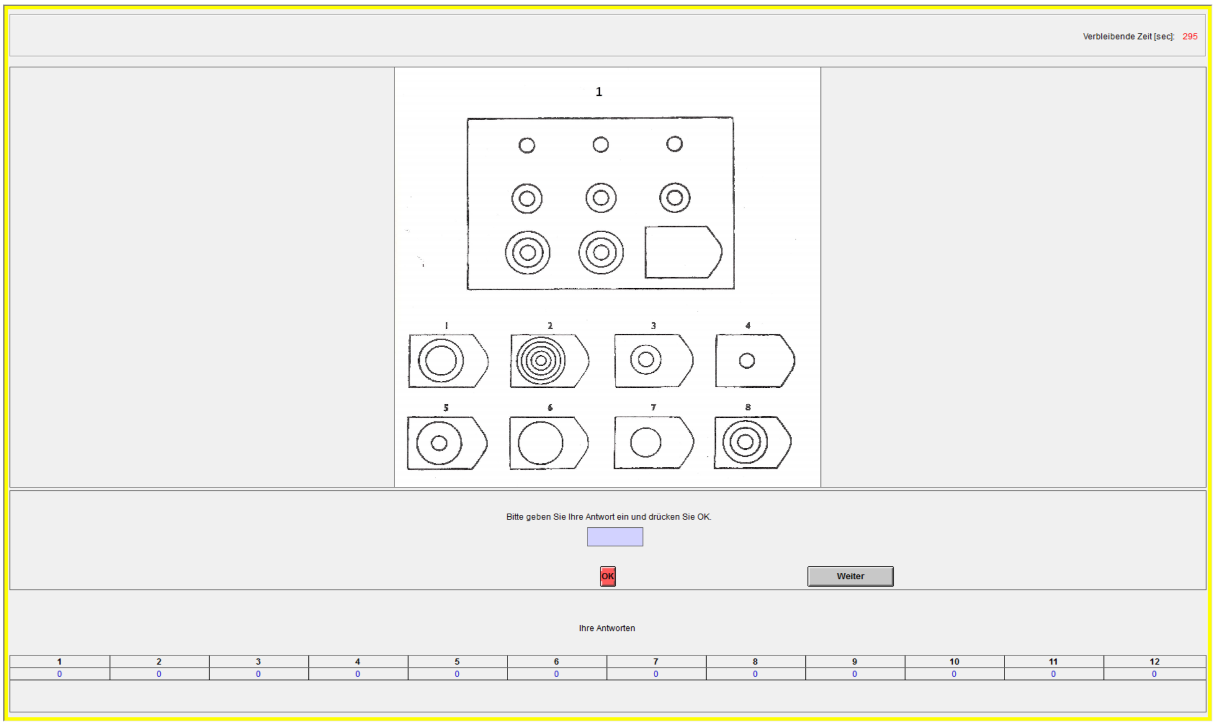

Figure 11: Raven Test Screen

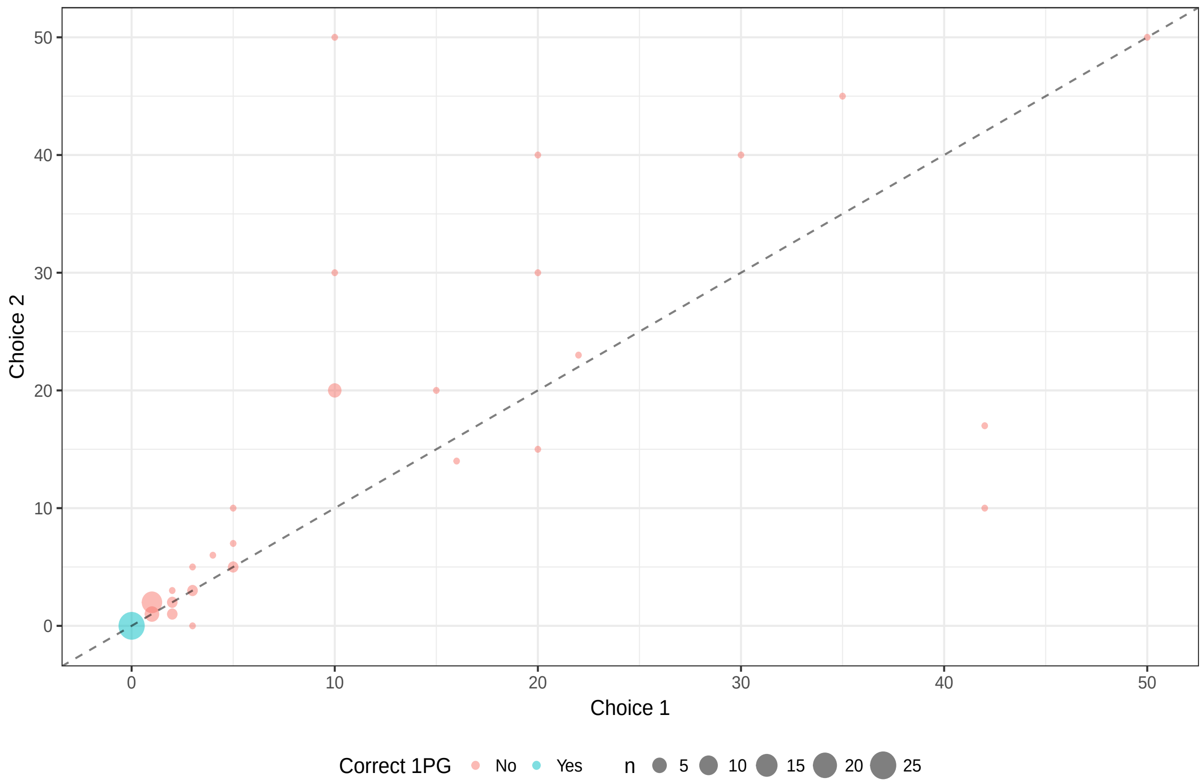

Figure 12: Zoom-in depicting only choices [0, 50] of Figure 1. In light blue the choices made by those that fully solved the 1PG.

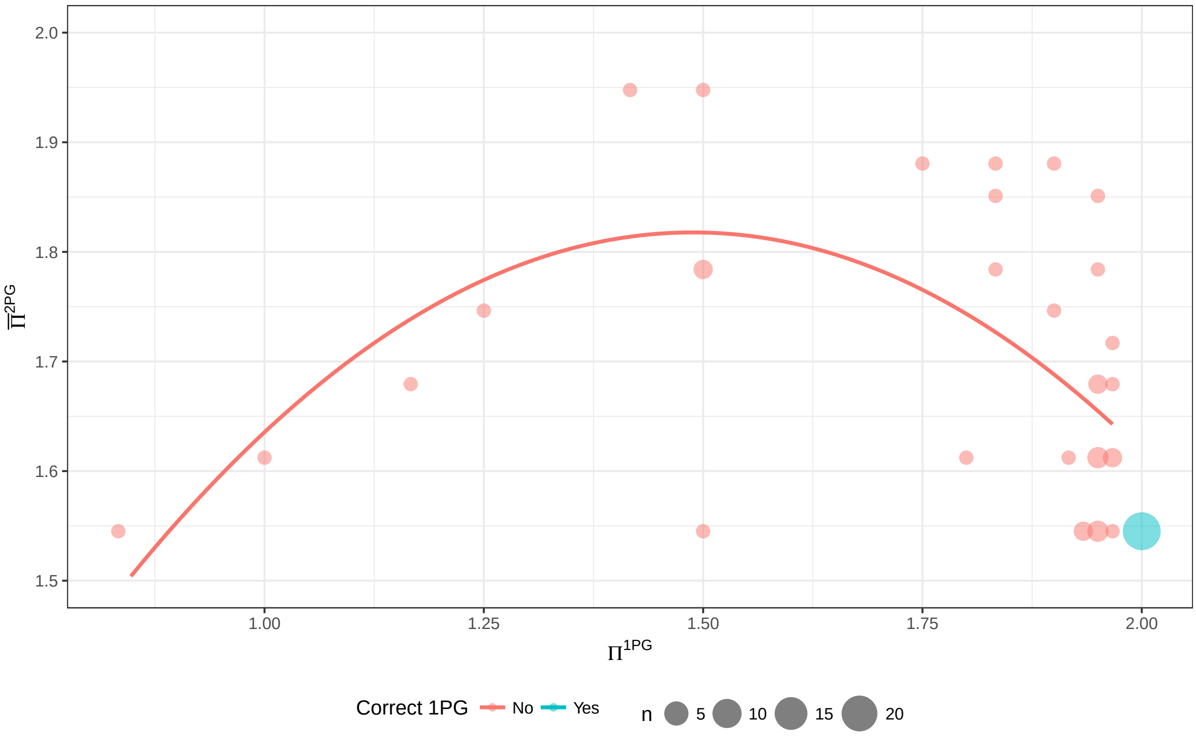

Figure 13: Zoom-in depicting payoffs [1.5, 2] of Figure 4. In light blue the choices made by those that fully solved the 1PG.

Figure 14: Plot with.

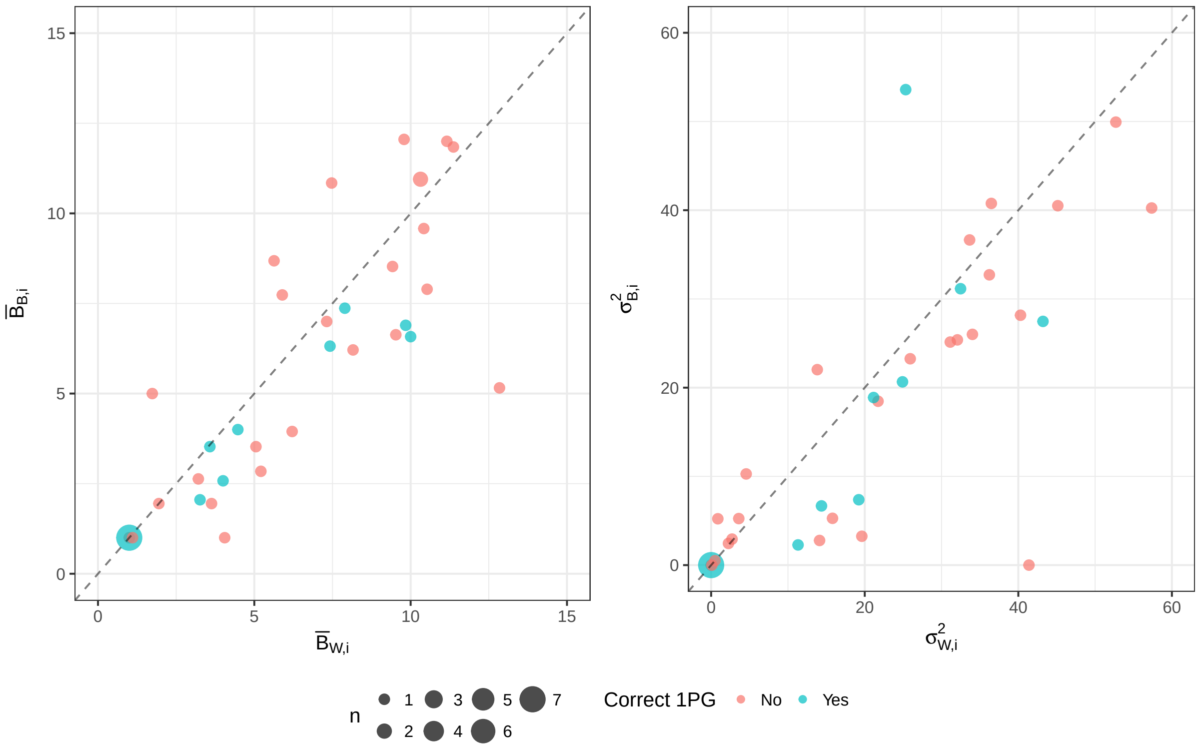

Figure 15: Figure depicting the change in distribution across the original belief elicitation task and WI. In the left(right) panel on the vertical axis the average value(variance) of the original distribution for subject i,in the horizontal axis the average value(variance) for the WI distribution for subject i. In both cases the dotted line is a 45 degree line; any value to the right of it represents an increase in the parameter in the WI distribution with respect to the original distribution.

D Instructions

The instructions below are translated from the original German instructions. The instructions were distributed sequentially, each of the following subsections at a time. After reading instructions, subjects completed the tasks, and then received instructions for the following task. Subjects were given time to carefully read the instructions and ask questions.

D.1 Overview

Today's experiment consists of a sequence of questionnaires and games. After each game or questionnaire, you receive instructions for the next game or questionnaire. In total, the experiment will last about 1.5 hours. For your participation your will receive a show-up fee of 5€. Depending on how you answer/play the questionnaires/games you can earn money on top of that. After you have finished all questionnaires and games, your final payoff will be shown on your screen. You will then receive a receipt, which you please fill out with the shown payoff and your name and address. After this you will be called to the adjoining room to receive your payment.

You will now receive the specific instructions for the first game.

D.2 Task 1

In this task, you have to pick two numbers (named X and Y) between 0 and 100 (both inclusive). Your payment for this task depends on how close these two numbers (X and Y) are to the so called "target number". The closer your numbers are to the target number, the higher your payment.

Caclulation of the target number The target number depends on your picked numbers. It is calculated as the mean of both numbers (X and Y), multiplied by two thirds:

$$\text{target number} = \frac{2}{3}\left(\frac{X+Y}{2}\right)$$

Your payment Your Payment depends on absolute distance of your numbers to the target number. For each number, if the number is exactly equal to the target number, you receive 1€. Should you not hit the target number exactly, some money will be deducted.

For each absolute unit distance of your numbers to the target number, 0.05€ will be deducted:

$$\text{payment for X} = 1€ - 0.05€\,|X - \text{target number}|$$

$$\text{payment for Y} = 1€ - 0.05€\,|Y - \text{target number}|$$

Because only the absolute difference between your number and the target number is used to calculate your payment (indicated with by the vertical bars "|"), it does not matter whether you deviate upwards or downwards - only the absolute difference counts. Please note that your payment can not become negative. Should your payment according to the above formula become negative, you will receive 0€ instead.

You have four minutes time to pick both numbers. After you have chosen your numbers, please confirm your choice by clicking the red "continue" button. Otherwise, your answer will not be recorded.

D.3 Task 2

This task is similar to the previous task. This time, however, you will play against another person in this room, who will be matched with you randomly. This time, you choose only one number between 0 and 100 (both inclusive). The other person also chooses a number between 0 and 100. Your payment now depends on how close your number is to the so called "target number". The closer your number is to the target number, the higher your payment.

Caclulation of the target number The target number depends your number, and the number stated by the other person. It is calculated as the mean of both numbers, multiplied by two thirds:

$$\text{target number} = \frac{2}{3}\left(\frac{\text{your number} + \text{other player's number}}{2}\right)$$

Your payment Your Payment depends on absolute distance of your numbers to the target number. If your number is exactly equal to the target number, you receive 2€.

Should you not hit the target number exactly, some money will be deducted. For each absolute unit distance of your number to the target number, 0.1€ will be deducted:

$$\text{payment} = 2€ - 1€\,|\text{your number} - \text{target number}|$$

Because only the absolute difference between your number and the target number is used to calculate your payment (indicated with by the vertical bars "|"), it does not matter whether you deviate upwards or downwards - only the absolute difference counts. Please note that your payment can not become negative. Should your payment according to the above formula become negative, you will receive 0€ instead.

Since the payoff of the other player is calculated in the same way, s/he also has the incentive to state the closest number possible to the target number.

You have four minutes time to pick your number. After you have chosen your number, please confirm your choice by clicking the red "continue" button. Otherwise, your answer will not be recorded.

D.4 Task 3

Reminder: in the previous task every participant was asked to pick a number between 0 and 10 (both inclusive). The payment for this game depended on how close this number was from the target number (mean of the numbers of both players, multiplied by two thirds).

In this task we would like to ask you to estimate how the other 19 participants who are participating in this experiment, behaved in the previous game. For this purpose, you receive 19 "token", that you can distribute in a number of bins (see screenshot below [Figure 10]).

Each token represents your estimate for one participant. The bins represent an interval for your estimation. To distribute tokens into bins, you type the number of participants, who you estimate to play a number in the interval of the respective bin, into the respective field on your screen. If you believe that no participant played a number in any interval, you can leave the field empty (or type in 0).

Example: You believe that 5 participants chose numbers between 75 and 79, and 14 participants chose numbers between 20 and 24. In this case you type "14" in the field 20-24 and "5" in the field 75-79.

For each token that coincides with a decision of a participant, you receive 0.20€.

Example: You have put 5 tokens in a bin, and 2 participants have actually played numbers in this bin. In this case, you receive 2*0.20€. If 5 people have actually played numbers in this bin, you receive 5*0.20€. If you estimate all 19 participants correctly, you receive 19*0.2€=3.80€.

Note: because it is difficult to keep track of how many tokens you have already distributed, we have built in a "check-button", that sums up all distributed tokens. Should not all 19 tokens be distributed, or more 19 tokens distributed, a pop up will tell you this. If you are happy with your distribution, and have distributed 19 token, you can finish this task by clicking "Ok".

D.5 Task 4

[Only a subset of 40 subjects participated in this task]

A few weeks ago, we conducted an experiment with 20 participants, who also solved Task 2. The participants received the exact same instructions as you, however, they have not solved Task 1 before they solved Task 2. That means, they played directly against another player, without picking two numbers before in Task 1.

We would like to ask you, similarly to the previous task, to estimate how these participants have played in Task 2. Again you receive 19 token, that you can distribute in a number of bins.

Since 20 participants took part in this experiment, but you have 19 tokens, we will randomly choose 19 participants from these 20 participants.

Again, for each correctly placed token, you receive 0.20€.

D.6 Task 5

In this task, you see a puzzle on every screen: a matrix containing 8 graphical entries and one empty field. There are 8 given option to complete the puzzle - only one is correct. For each puzzle, you have to select one of the 8 options. To this, you type in the corresponding number, and confirm by clicking OK.

[screenshot of Figure 11]

In total you will complete 36 such puzzles, divided into 3 blocks of 12 puzzles. Within each block you can move back and forth between puzzles, and correct previously stated answers. You have 5 minutes to solve each of the first two blocks, and 8 minutes for the remaining block. The remaining time will be displayed in the header of your screen. After you have completed all puzzles, or the time runs out, you have to confirm your answers by clicking "send answers". If you do not confirm your choices, they will not be recorded.

For each correctly answered puzzle, you earn 0.10€.

D.7 Task 6

This task consist of the game "Race to 60", that you play repeatedly against the computer. Your task is to win this game as often as possible against the computer.

In this game, you and the computer player state numbers between 1 and 10 sequentially. The numbers are added up, and whoever is first to push the sum of all stated numbers to or above 60 wins the game.

The game functions as follows: You start the game by stating a number between 1 and 10 (both inclusive). Then the game continues as follows:

1. The computer chooses a number between 1 and 10. This number wil be added to your number.

2. The screen shows the sum of all numbers chosen so far. If the sum is smaller than 60, you again choose a number between 1 and 10, which again will be added to all previously stated numbers.

This will be repeated until the sum of all numbers is greater or equal to 60. Whoever states the number that pushes the sum to or above 60, wins the game. You will play this game 6 times against the computer, and you have 90 seconds time for each repetition. For each game won, you receive 0.50€.

Your payment This is the last game for today's experiment. After completing this game, you will receive a questionnaire that you please complete. After that, your payments will be summarized on the screen. You then receive a receipt, that you please fill out. After you have done that, an experimenter will come to your place to check if everything is in order. After that you can go to the adjoining room to receive your payment.