

Hoffmann, Timo

Conference Paper

The Effect of Belief Elicitation Game Play

Beiträge zur Jahrestagung des Vereins für Socialpolitik 2014: Evidenzbasierte Wirtschaftspolitik
- Session: Experimental Economics II, No. B09-V1

Provided in Cooperation with:

Verein für Socialpolitik / German Economic Association

Suggested Citation: Hoffmann, Timo (2014) : The Effect of Belief Elicitation Game Play, Beiträge zur Jahrestagung des Vereins für Socialpolitik 2014: Evidenzbasierte Wirtschaftspolitik - Session: Experimental Economics II, No. B09-V1, ZBW - Deutsche Zentralbibliothek für Wirtschaftswissenschaften, Leibniz-Informationszentrum Wirtschaft, Kiel und Hamburg

This Version is available at:

<https://hdl.handle.net/10419/100483>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

The Effect of Belief Elicitation on Game Play

Timo Hoffmann ^{*}, [†]

February, 2014

WORKING PAPER - COMMENTS ARE WELCOME

Abstract

The assumptions that subjects hold beliefs and that the chosen actions are not altered by a proper elicitation of these beliefs are widely used in economics. In this paper I experimentally test whether the second assumption is correct. Especially controlling for different game properties, I find that in dominance solvable two-player normal-form games belief elicitation results in a significant increase of equilibrium play. Therefore the elicitation of beliefs can affect the choices made by subjects and lead to more equilibrium actions being chosen. Surprisingly one major reason for this effect is the decreased play of own dominated actions. The results indicate that belief elicitation induces subjects to “think harder” about the presented decision situation, which results in a better understanding of the given situation and consequently in a modification of their beliefs. Therefore, in certain decision situations, belief elicitation affects the decisions made by subjects.

JEL classification: C72, C91, D83

Keywords: Experimental economics; Elicitation; Stated beliefs; Normal-form games

1 Introduction

Understanding, describing and explaining the behavior and choices of individuals is one of the core objectives of economics. Observed choices are mainly used to infer underlying rules and mechanisms, which drive subjects' behavior, and to formulate models, which explain and predict the observed choices. While observed choices are useful to infer these rules, they are not always sufficient to distinguish between competing theories and explanations. The game-theoretical models that are used generally make the assumption that all players hold beliefs about the actions other player(s) will chose. In a two-player simultaneous strategic game beliefs of a player describe his or her expectation about the likely behavior of the other player. One of the earlier papers that elicited beliefs was McKelvey and Page (1990) in the

^{*}I appreciate the comments and the advice received from Dirk Engelmann, Werner Güth, Henrik Orzen, Stefan Penczynski, Philipp Schmidt-Dengler and Roberto Weber. I thank the participants at various seminars in Mannheim, as well as the audiences at the 2013 Workshop on Behavioral and Experimental Economics in Florence, the 2013 ESA World Meetings in Zurich, the 2013 GfeW conference, and the 2013 EEA conference and the 2013 Workshop on Experimental Economics in Xiamen for many helpful comments and fruitful discussions. I also want to thank Carsten Schmidt for his help in setting up the experimental sessions. All remaining errors are my own.

[†]Corresponding address: Center for Doctoral Studies in Economics, University of Mannheim L9, 7, D-68161 Mannheim, Germany, e-mail: Timo.Hoffmann@gess.uni-mannheim.de

context of information aggregation. The beliefs in their experiment are elicited in a non-strategic setting and since then many studies have used belief elicitation. It has being widely used in the literature about fairness and reciprocity in which the beliefs about the giving or sharing behavior of the other subject(s) have been elicited (e.g. Offerman et al. (1996), Dufwenberg and Gneezy (2000), Gächter and Renner (2010)). Furthermore belief elicitation has been essential in the literature about (belief) learning models like Nyarko and Schotter (2002) (see Camerer (2003) chapter 6 for an extended overview), the literature about behavior in one-shot games and the probability of equilibrium play (e.g. see Costa-Gomes and Weizsäcker (2008), Costa-Gomes et al. (2001)) as well as in many other areas.¹ Besides the assumption that subjects have these beliefs, at least in theoretical models, it is usually assumed that the chosen actions of players are best responses to these beliefs. This is an assumption that has been tested in several papers (e.g. Huck and Weizsäcker (2002), Weizsäcker (2003), Costa-Gomes and Weizsäcker (2008)). If actions are based on beliefs, knowing beliefs helps to understand the reasoning and intentions of players for choosing a certain action and provides additional information which can be used to differentiate between competing models that predict or explain a similar choice pattern.² In particular beliefs might be very insightful to understand the low level of equilibrium play observed in many experiments (see Kagel and Roth (1995) and Camerer (2003)). Therefore stated (and also inferred) expectations of subjects have been widely used in order to determine the reasoning of subjects and to explain or classify their behavior.³ The frequent use of the elicitation directly gives rise to the question whether choices are affected by these elicitation. This paper wants to answer this question by comparing the choices of subjects in two different treatments. While in the first treatment subjects state their actions choices without being asked about beliefs, in a second treatment a separate group of subjects plays the exact same games, but the subjects are always asked to state their actions choices and their beliefs about the behavior of the other player simultaneously.

The question whether the elicitation of beliefs affects the action choices of subjects is important for multiple reasons, because if it is answered positively this has consequences for the comparability of results obtained in studies with and without belief elicitation. Many papers using belief elicitation simply assume that the behavior (and beliefs) of subjects are not altered by these procedures. Fischbacher and Gächter (2010) for example are concerned about possible ordering effects in their experiment and therefore conduct several treatments in which they exchange the order of their two experiments. But they elicit beliefs in all their treatments. Since the reported behavior in their public good game is similar to the behavior reported in related experiments they extend their findings also to public good experiments without belief elicitation. Doing this is only possible if the behavior is not affected by the elicitation of beliefs. Similarly in Offerman et al. (1996) there is no separate treatment in which the public good game

¹For additional references regarding papers with belief elicitation see the papers cited in the introduction of Costa-Gomes and Weizsäcker (2008) or Blanco et al. (2010).

²While beliefs are often used to derive best responses for subjects, it is not always the case the beliefs are assumed to be existent. Often, especially if they are not directly elicited, beliefs are also seen as a theoretical construct to explain or rationalize observed behavior.

³In a public good game with voluntary contributions Fischbacher and Gächter (2010) use a classification of subjects proposed by Fischbacher et al. (2001) to distinguish different player motives for cooperation. They sort subjects into different types according to their stated preferences and beliefs and are therefore able to distinguish between various motives for positive as well as for zero contributions. Using this classification they can show that the usually observed breakdown of cooperation in these public good games is not dependent on the presence of free riders, but can be explained by imperfectly conditional contributors.

was played without belief elicitation. Others authors acknowledge a possible effect of the elicitation on the stated choices (and beliefs), but argue that the observed behavior is similar to the behavior reported in closely related studies (e.g. Haruvy et al. (2007)) and therefore that additional treatments without the elicitation procedure are unnecessary. But it is not a priori obvious that the elicitation is without effect on subjects' decisions. Especially if the first assumption, that all subjects always form beliefs and therefore the elicitation neither induces nor changes beliefs, is violated, it is no longer given that the elicitation of beliefs has no influence on the chosen action. Also, if the choices are affected by the belief elicitation this finding calls into question the standard assumptions that agents hold beliefs and base their decisions on them. More specifically, if an effect can be identified the relevant question becomes through which channel the elicitation affects the choices. While such an effect creates some doubts regarding the assumptions that subjects always hold beliefs, the effect could also be triggered by an influence on the previously existing beliefs.

Not all papers simply assume that there is no effect of the elicitation, some papers explicitly test for the existence of a possible effect of the elicitation on observed choices. The results obtained by various studies are mixed. Probably the first studies which are concerned with the question whether belief elicitation changes game play is Croson (1999, 2000). She tests for an effect in a public good game with a voluntary contribution mechanism and a prisoner's dilemma game and reports that belief elicitation is not neutral, but influences which actions are chosen. In both games the theoretical predictions of no contribution (public good game) and defection (prisoner's dilemma) are dominant for all subjects. Croson (2000) finds that in treatments in which subjects additionally to their choice report a guess of the sum of the others' contributions (public good game) or the choice of the randomly assigned partner (prisoner's dilemma) the choices are more in line with the theoretical predictions. Subjects tend to contribute less and defect more often. Due to these findings Croson claims that the behavior of subjects changes when beliefs are elicited and she argues that belief elicitation leads to more equilibrium play.

Several studies have tried to replicate the findings of Croson regarding the effect of belief elicitation in the public good game. Their findings are mixed. While Wilcox and Feltovich (2000) do not find an effect, Gächter and Renner (2010) report that contributions actually increase in the belief elicitation treatment. Zelmer (2003) concludes from her meta-study of public good games that belief elicitation is linked to decreases in contribution levels, even though in her sample there is only a very limited number of experiments that used belief elicitation. Nyarko and Schotter (2002) in a repeated 2x2 game, mainly used for the estimation of belief learning models, also test for an effect of belief elicitation on game play. They report no differences between the treatments with and those without belief elicitation. Rutström and Wilcox (2009) are explicitly concerned with the possible effect of belief elicitation on the chosen actions of subjects. They use an asymmetric matching pennies game played repeatedly and they find that paid belief elicitation changes game play significantly for subjects with a large asymmetry of payoffs. The effect is larger in the first rounds of play. The experiment of Costa-Gomes and Weizsäcker (2008) analyzes equilibrium play in normal-form games and as part of their study they also test whether there exists an effect of belief elicitation on game play. They do not report any significant effects on the chosen actions.⁴ In a recent survey article about belief elicitation Schlag et al. (2013) conclude that there is a

⁴While they check for an effect of belief elicitation there are two notable differences of their analysis to mine. First, Costa-Gomes and Weizsäcker have a treatment in which actions and beliefs are stated for each game, but the elicitation is always done

relatively limited number of studies on the effect of belief elicitation on choices and that the results from these studies are inconclusive. Summarizing this literature there are two different facts that emerge. Most importantly belief elicitation does not always influence game play. But also if studies find an effect the elicitation of beliefs does not always drive subjects towards the equilibrium, contrary to the observations in Croson (1999, 2000).

The design of the experiment conducted in this paper is such, that not only differences in subjects' choices with and without belief elicitation can be detected, but it further allows me to test if different properties of the strategic situations have an influence on the existence (or strength) of such an effect. I believe this to be important, since belief elicitation might affect choices in various situations differently. Therefore this paper does not only add to the question whether belief elicitation can have an effect at all, but allows for a more detailed analysis under which conditions the effect is most likely to occur. This is an important extension of the previous existing work in this area. Hence the results obtained with my experiment can also help to advance our understanding of why belief elicitation might be relevant for the choices of subjects. Further, since the difference between the treatments is the timing when the beliefs of subjects are elicited (simultaneous with the action choice or as a surprise after all decisions have been made) the design provides an opportunity to compare the stated beliefs of all subjects as the expectations are elicited in all treatments. This is not the case in most other papers, which only elicit beliefs in some experimental conditions. Due to the belief elicitation in both cases it is possible to compare the stated beliefs in the two treatments and therefore to test, whether the difference in the timing affects the chosen actions or the stated beliefs.

In contrast to the leading hypothesis that belief elicitation does not have a significant influence on the behavior of subjects, I find that in the treatment in which subjects state their expectations, they on average chose more actions in accordance with Nash equilibrium (54.6 percent in the treatment in which choices are stated without beliefs and 58 percent in the treatment, in which choices and beliefs are stated simultaneously). The difference is most pronounced and highly significant for games that are dominance solvable (60.8 percent compared to 68.4 percent). Surprisingly, one main driving force for the effect seems to be an increased recognition and avoidance of own dominated actions, rather than a reaction to the existence of dominated actions of the other player. While there exists an effect on the chosen actions, there seems to be no difference between the beliefs stated simultaneously with the decision and those stated after all decisions are made. This indicates that the elicitation directly affects beliefs and that the difference in actions are a result of these modified beliefs, a finding which supports the assumption of belief based play. Hence, in this paper I provide evidence for the claim that belief elicitation results in higher shares of equilibrium play. As I will argue in more detail below there are indications that suggest that the effect is triggered by subjects that "think harder" when beliefs are elicited.

This paper is organized as follows. The next section states two hypotheses why game play might be influenced by belief elicitation. Section 3 describes the experiment and procedures in detail, while section 4 provides the results and analyses and section 5 concludes.

before the action choice is reported. In my design the two tasks are done simultaneously. Second, while they find no differences on the game level they do not analyze the individual behavior across games explicitly as I do in section 4.1.

2 Hypotheses

Most game theoretic concepts like Nash equilibrium or somewhat weaker concepts like rationalizability (Pearce, 1984; Bernheim, 1984) assume that players have (subjective) beliefs, meaning they have an idea in mind what the other player will do. Often times how these beliefs are derived, e.g. through introspection or experience, is not specified and without relevance for game play. While different concepts differ in their requirements on these beliefs, it is a common feature of them that each player has these beliefs. Therefore an elicitation of these beliefs should theoretically not change the behavior compared to a situation without belief elicitation, since a simple statement of something that players have anyway does not change the game or the optimal behavior. The result predicted by these theories might be slightly different if subjects are paid to state correct beliefs, e.g. if a (proper) scoring rule, like the quadratic scoring rule, is used. The additional payoff can change the incentives to play the game and therefore lead to a difference in behavior, especially if subjects are not risk neutral.⁵ In order to avoid this problem, my experimental design takes care of possible hedging opportunities, by only paying the belief statement or the chosen action for a given game. Therefore subjects cannot use the belief statement to insure themselves against unwanted outcomes. Furthermore, while I rewarded the stated expectations of subjects the procedure does not increase the expected payoffs of a single game. Because only the belief or the outcome of a game is paid and the payoff from the belief statement is usually smaller than the payoff for the outcome, the additional payment does not increase the payment for the outcome of a single game. Therefore the action choices in all treatments are paid equally. This is important if one believes that an increased payment would lead to different behavior. But since the main payoff for subjects comes from the chosen action the reward for the elicitation does not increase the amount “at stake” for an individual game. Given that in all treatments subjects know that only the result of their chosen action or the stated belief can be paid in the same game and since the likelihood that the outcome is paid is identical in all treatments, the monetary incentives the subjects face are always equivalent. Hence there is no reason to expect a change in the chosen actions due to the payments made for accuracy of the stated beliefs.

There are two possibilities how the elicitation of beliefs might affect the choices made by subjects in one-shot games.⁶ On the one hand it could be that play is not belief based, such that contrary to the standard assumption, players do not hold expectations about the likely behavior of the other player or that they do not take these beliefs into account when making their own action choice. On the other hand the elicitation might not lead to the formation of beliefs, but affect these beliefs by increasing the awareness for the strategic character of the game or leading to a more involved analysis of the likely behavior of the other player. If subjects do not have beliefs about the behavior of the other player it could be that subjects use other mechanisms or (simple) heuristics to determine which action to choose. A subject might always choose the action that yields him or her the highest possible payoff or select the action, which ensures him or her the highest payoff for sure, independent of the choice of the other player (max-min choice). It could also be that subjects choose their action completely randomly, but this is something that is

⁵Blanco et al. (2010) conduct an experimental analysis of the effect of incentivized belief elicitation and the possibility to use incentivized belief statements as a hedging device. A general message to take away from their paper is that one has to be careful if significant hedging opportunities arise due to the incentivized belief elicitation.

⁶In repeated games the question about belief learning and experience also becomes relevant, but since in this study I concentrate on one-shot environments these aspects are not relevant.

clearly rejected in the data. Such strategies do not require subjects to form any expectations about the behavior of the other player, since their own choice does not depend on it. While such behavior might be unlikely in situations in which subjects have a lot of experience or which are repeatedly encountered it is more likely in the studied context of one-shot games. In those situations the question whether subjects have meaningful beliefs at all arises. If the explicit request to state beliefs leads to the formation of beliefs, there is no reason to expect that the resulting actions are identical to a situation without belief elicitation. Since a subject that does not form a belief without the explicit elicitation might not have these expectations for different reasons. It could be that a subject does not consider the behavior of the other player relevant or that he or she has no idea what the other player will do. While it is highly unlikely that no subject considers the action of the other player when deciding which action to choose in all games, for the belief elicitation to significantly influence the choice behavior a certain fraction of subjects without beliefs could be sufficient. Also some subjects might only form beliefs for some, but not all games or only realize that they should consider the likely action of the other player after having already stated their decisions for several games.⁷ When there exist at least a few subjects that do not take the likely behavior of the other player(s) into account the elicitation procedure should induce some beliefs for these subjects and therefore the overall distribution of actions can be significantly altered. As they are required to form beliefs, it is to expect that they use these expectations to derive their decisions. This would make it very likely that the actions chosen by these subjects are different in a situation with belief elicitation compared to choices made without this elicitation. Hence, if the assumption that all subjects hold beliefs is false, I should observe significant differences in the choice patterns with and without belief elicitation. Such an effect should be present for all games, but will most likely be more pronounced in games in which the formation of beliefs has a decisive effect on their own decision. Given that beliefs are formulated and subjects then best respond to these beliefs it is more likely that a player chooses an action that is part of a Nash equilibrium. The observed difference should be larger in games that have a unique equilibrium or are dominance solvable, since in these situations the recognition of dominated actions of the other player should be a likely consequence of the formation of (reasonable) beliefs. These detections have direct consequences for players' own action choice in such games. As a result actions in accordance with equilibrium play are expected to be more frequent in all games in sessions with belief elicitation, if the elicitation leads to the formation of beliefs. The effect should be larger in dominance solvable games.

Given the strategic character of the games used in this paper, I consider it unlikely that many subjects do not form any expectations about other subjects' behavior when they are not explicitly asked to state them. Also in the instructions of the experiment and in the explanation of the possible payoffs for the games it is made clear that the outcome of a game is determined by the decision of both players. It is clearly communicated that the action choice of the other player matters for the own payoff. Therefore I expect subjects to take the likely action of the other player into account when choosing which action to take themselves even when they are not explicitly asked for them. So the elicitation of these beliefs should not be the reason for them to have beliefs at all. But it could still be the case that the elicitation has an influence on the formulation of these beliefs. While subjects might have a rough expectation about what the other player will do, the request to state their beliefs explicitly could induce a more involved

⁷Although such a learning could take place, the fact that no feedback about the choices of the other players are given until the very end of the experiment makes only "no feedback learning" possible, which is not very likely.

thought process. Without being asked to put exact numbers on the expected behavior of the other player a subject might have an idea what is “more likely to happen”, but the additional task might motivate subjects to think more about the game in general and about the choice of the other subject in particular. In order to answer the question how likely is it that the other player chooses a certain action exactly the player might increase his or her effort and put him- or herself “in the shoes of the other player”. This increased effort could lead to a modification of the “original” beliefs (the beliefs a player has without being explicitly asked to state them), since the increased involvement with the game ad hand might result in a deeper understanding of the decision situation and therefore lead to a modification of beliefs.

Different to the hypothesis outlined above subjects are not forming beliefs just because they are elicited, they already consider the likely choice of the other player in their decision, but they might revise their expectations because of the required explicit belief statement. Such a refinement of their beliefs is most likely to change behavior if the resulting expectations have a direct consequence for the own decision. Given that subjects hold beliefs even when they are not elicited it is likely that a significant effect on game play only occurs when the deeper reasoning yields a decisive shift in expectations. A change in beliefs is important for behavior if this belief modification results from some knowledge about an important aspect of the situation which was not observed before. Therefore an effect of the elicitation on all games is less likely, since not in all games deeper thinking leads to the discovery of an important aspect that influences which action a player chooses. But the effect of the belief elicitation could be strong for dominance solvable games. In these games the other player usually has an action that is dominated and therefore should not be chosen. There are two kinds of dominated actions and both are more likely to be detected when subjects think more about the game. An action of the other player might be dominated “from the start on” meaning independently of the choice of his or her opponent. Belief elicitation should not influence that a player realizes that he or she him- or herself has such a dominated action, but it should affect the likelihood that the other player detects this domination as well. Given this observation a player should place a zero (or very small probability) on the likelihood that this dominated action is played. In dominant solvable games this then leads to an action of the player that is iteratively dominated. An iteratively dominated action is not dominated if a player puts some probability on all possible actions of the other player, but only once one action of the other player is ruled out this action becomes dominated and therefore should not be played. Since this is the case in dominance solvable games “thinking harder” should lead to an increased detection (and according) reaction of actions that are iteratively dominated. With a deeper understanding of the situation the subjects should recognize dominated actions of the other player more often, which results in less frequent play of own iterated dominated actions. Hence I expect an increase in the number of actions chosen in accordance with equilibrium predictions for these games. Since this chain of reasoning is not present in games that are not dominance solvable belief elicitation might not necessarily result in different behavior for other (not dominance solvable) games.

This “thinking harder” hypothesis explains an effect of the belief elicitation on the chosen actions with an increased effort and subsequent better understanding of the game by subjects. The additional task has subjects think more about the game and therefore they behave “more like a game theorist” (Croson, 2000). This explanation predicts more equilibrium play and can therefore explain the findings by Croson (1999, 2000). But belief elicitation cannot only induce more equilibrium play. It is also an additional task that subjects have to understand and “solve”. Instead of simply stating their chosen action

subjects have to understand how to state their beliefs and how they are rewarded for it. Therefore it should foremost increase the effort required to complete the task. The elicitation could result in more effort spent to understand the decision situation (with the expected results on behavior just outlined) or it leads to more confusion. If it leads to more confused subjects the elicitation could result in more random play. As a consequence the chosen actions might simply be different because subjects are more confused. Subjects could be irritated or simply overwhelmed with the additional task. If that is the case I would not consider it as an effect of the elicitation, but rather the result of poor instructions. Since all decisions in this experiment are made without any time pressure and subjects had to answer a series of understanding questions regarding payoffs and belief statements before the start of the actual experiment, I consider it unlikely that the additional task confused or distracted subjects. Nevertheless, if it would be the case it should lead to less equilibrium play, since there is no reason why a confused subject should play as predicted by a theory that has strong demands on subjects' rationality and strategic sophistication. Therefore confusion or distraction does not explain an increased play of actions in accordance with equilibrium and because of this belief elicitation would work in the opposite direction as predicted by the two hypotheses about the effect of belief elicitation on game play above.

3 Experimental design and procedures

3.1 Experimental design

Subjects played a series of 20 normal-form games with at least one pure-strategy Nash equilibrium. There are two treatments. In the baseline treatment subjects first play all 20 games without feedback. After they have completed all games they are asked to state their beliefs about the likely chosen actions by their respective opponent for each game. While doing this they see the games in the same order as before. This procedure ensures that they are informed about all payoffs when they state their beliefs, but they do not have any information about the behavior of the other player. They are also not reminded of their own choice in these games. The chosen actions in this treatment are compared to the behavior of subjects in the belief treatment. In this treatment subjects play the exact same 20 games, in the same order, but always state their beliefs and their decision for a game at the same time. Therefore subjects in this treatment are asked to indicate their choice and to state their beliefs simultaneously, while in the baseline treatment the action choice is stated without belief elicitation. Subjects could move from one game to the next at their own speed in each treatment, but once they submitted an action or belief they could never go back to previous games. After all subjects have completed all 20 games they received feedback about the outcomes in the paid games and were shown their payoff. Each subject only participated in one treatment.

All games were two-person games with two or three actions for each individual. This means the games were of the 2x2, 3x3 or 2x3 form.⁸ The games vary in multiple dimensions, explained in detail below, and subjects in both treatments played the same games. Of these 20 games four are symmetric, so the game is the same for the row and the column player. The other 16 games are actually eight distinct games, but subjects play each game twice, one time in each role. Subjects do not play the exact

⁸All games are shown in the Appendix.

game in both roles, but rather isomorphic transformations. Therefore there are eight pairs of equivalent games. For each asymmetric game a second game is created by transposing the player roles, changing the order of actions and adding or subtracting a constant to all payoffs. This leaves all relevant properties of the game (like the Nash equilibria and dominance relations) unchanged.⁹ Because of these non trivial transformations, it is unlikely that subjects realize that they play each game in both roles as the payoff and action order change disguises the equivalence of the games. Additionally two equivalent games are never played in direct sequence.

Important characteristics that differ between groups of equivalent games are whether or not a game is dominance solvable, if it has a unique or multiple pure strategy Nash equilibria, whether the Nash equilibrium is Pareto dominated by another outcome and whether or not there exists an alternative outcome to the Nash equilibrium with symmetric payoffs for both subjects. Also dominance solvable games differ in the number of rounds which are needed to arrive at the dominance solvable solution. The order in which the games are played is the same for all subjects in both treatments. Subjects in each treatment are split into column and row players. While the player roles are fixed over the treatment all games are presented as each player being a row player (column players are shown the transposed game). Due to the isomorphic transformations behavior of row players in asymmetric games can directly be compared to behavior of column players in the equivalent game and vice versa. Table 1 shows the properties for all 12 unique games.

Table 1: Game Properties

Overview of Properties for all Games												
Characteristic / Game #	1	2	3	4	5	7	9	11	13	15	17	19
unique NE		X	X	X	X	X	X	X		X	X	X
NE Pareto dominated	(X)	X	X		X	X				X		
symmetric game	X	X	X	X								
3x3	X		X	X		X	X	X	X			
3x2										X	X	X
2x2		X			X							
dominance solvable	X	X	X		X	X			X	X	X	
elimination of weakly dom. strategy	X								X			
column dominated action	X	X	X			X		X	X		X	
row dominated action	X	X	X		X					X		
symmetric alternative	-	-	-	-			X		(X)			X

Notes: Game numbers given in the table correspond to each game in the appendix. Of course there are symmetric alternatives in symmetric games. (X) for NE Pareto dominated means that the other outcome is also a NE, (X) for symmetric alternative means that no player is better off with the symmetric alternative compared to the NE.

⁹Games 5 and 6 for example are equivalent games. Game 6 is created by exchanging player roles (e.g. action T of the row player in game 5 becomes action L of the column player in game 6) and subtracting four points from all payoffs. The identical technique is used by Costa-Gomes and Weizsäcker (2008).

3.2 Experimental procedure

The experiment was carried out in the experimental laboratory of the University of Mannheim (mLab) in November / December of 2012 and March of 2013. Subjects were recruited from the subject pool of the University of Mannheim via e-mails using the ORSEE recruiting software (Greiner, 2004). The experiment was computerized and used the zTree software (Fischbacher, 2007).

Six sessions were conducted (three sessions for each treatment) with a total of 122 subjects, mostly undergraduate students of the University of Mannheim. In each session half of the subjects were randomly classified as “row” and the other half as “column”-players. All games were presented to the subjects as row players, hence no subject was aware of his or her role.

After reading the instructions all subjects had to pass an understanding test and only after successful completion proceeded to play all 20 games. In the baseline treatment subjects first entered their choices for all 20 games and only after the completion of this first part they were informed about the content of the second part. They were given instructions for the second part after all subjects had entered their choices for all games and had to pass a second understanding test, before they continued to state their beliefs. In this way it was ensured that the actions were chosen before the subjects were asked about beliefs. When subjects were asked for their beliefs they always saw the payoff matrix of the game. Beliefs can be stated as any two / three numbers adding up to 100. Subjects were not required to enter integers. Therefore the elicitation of beliefs elicited the probability distribution over all possible actions of the other player. The belief treatment consisted only of one part, since actions and beliefs were stated at the same time.

For each game subjects were matched with one player of the other type. This matching was changed after each round and completely anonymous. Subjects received no feedback while stating their actions and beliefs. Only after all subjects had entered all their beliefs and actions they were informed about their payoffs and the chosen actions of the respective opponents.

Subjects were paid for six randomly chosen games. For three games the payoffs resulting from the chosen actions were paid with subjects receiving four cents for each point. When choosing their action subjects always know that the outcome in three out of the 20 games would be paid. Also for three different games the belief statements were paid using a quadratic scoring rule to reward subjects for more accurate beliefs. Subjects were informed that their expected payoff from their stated beliefs are maximal if they state their true beliefs.¹⁰ Given the chosen parametrization of the quadratic scoring rule, which limits the maximal payoff for be stated beliefs to 2€, the payoffs for the actions were usually higher than for the stated beliefs. Also the probability that the outcome in an individual game was paid out was identical in both treatments. This and the fact that never the belief statement and the outcome of the same game were paid out ensures that the incentives for the action choices are identical across both treatments. The sum of all six paid games was the payoff for a subject. On average subjects earned 15.24€ and a session lasted always less than 90 minutes.¹¹

¹⁰The exact formula for the payment of the belief statements is given in the appendix. For each belief statement subjects earned between 0 - 2€.

¹¹This average contains the payoff of only five sessions, since in a sixth session due to a computer crash, that occurred after all decision and belief statements had been made, it is impossible to recover the matching. Therefore subjects received a fixed payment. The datafile containing all decisions and belief statements could later be recovered and the data is used in the

4 Results

4.1 Differences in chosen actions

To begin the analysis I test whether there exists a significant difference in the actions chosen between types for equivalent games. This means I compare the behavior of type-I (row players) and type-II (column players) subjects within each treatment. For the asymmetric games, I compare type-I subjects with the choices of type-II subjects in the transformed game (and vice versa). Such a test compares categorical data between independent samples, therefore I use two-sided Fisher's Exact Probability Tests (FEPT) for count data. In each treatment only one game has a significantly different distribution of actions between the two types (five percent significance level).¹² These two rejections in 40 comparisons are in line with the expected number of rejections. This result allows me to pool the data from the isomorphic games for each treatment for the following analysis, which greatly improves the power of all tests.

In order to get an indication for the power of the used tests, I analyze whether the chosen actions are different from random play. To evaluate this I perform χ^2 goodness of fit tests separately for each game. Testing the chosen actions yields a non-significant difference from the uniform distribution in six cases in the baseline and for two games in the belief treatment. These 32 rejections for 40 comparisons are clearly more than random rejections and therefore the tests show that in each treatment the actions chosen by the subjects are frequently and significantly different from randomness. Pooling all action choices over the two treatments and performing the same test on the pooled data leads only to two non-significant differences. Summarizing the results so far one can say that player roles are without effect on game play and that chosen actions are not random.¹³

Next I test, whether there is a treatment effect regarding the chosen actions. To test this I again perform FEPT, but compare the chosen actions between treatments. Testing all 20 individual games for differences between the treatments, I find a significant difference only in two games. While this number of rejections is slightly above the level of chance, it does not imply a strong treatment effect.¹⁴ While the difference is not significant in most games, it is interesting to note that in eight out of the ten games that have a unique NE and are dominance solvable, the share of NE actions chosen is higher in the belief treatment. Nevertheless, I cannot identify a treatment effect on the individual games.

Since looking at the data of the individual games does not reveal a lot of information about what is happening between the treatments, I use several aggregate measures to identify a treatment effect. To answer the main research question whether the elicitation of beliefs has an influence on the decision of subjects, I calculate for each subject the total number of actions chosen in accordance with NE predictions. The measure “# NE actions chosen” counts how often an individual chooses an action that belongs

following analyses.

¹²The p-values of game 11 in the baseline treatment and game 8 in the belief treatment are below 5%. All other p-values are above the five percent level and often are much higher.

¹³The games 2, 5, 6, 11, 15 and 16 are not significantly different from random play in the baseline treatment and in the belief treatment only for games 11 and 15 no significant difference can be found. In the pooled data again only the choices in games 11 and 15 are not different from random play.

¹⁴The difference is significant in games 16 and 18. The result is unchanged if not all actions are compared, but only whether an action that is part of a pure-Nash equilibrium was chosen or not. In both cases the significance is at the 1% level. It is also interesting to note that in both cases the frequency of equilibrium play is larger in the belief treatment.

to a pure strategy NE of the game. In order to determine a possible treatment effect the mean and the distribution of this measure are compared between the two treatments.

Overall, subjects in the baseline treatment play an action in accordance with NE in 54.6 percent of all cases. This is in line with other results on normal-form games.¹⁵ The share of equilibrium actions chosen increases to 58 percent in the belief treatment. Therefore, compared to the baseline treatment more subjects play NE actions. With the sample size of 122 subjects (60 in baseline and 62 in the belief treatment) this difference is not statistically significant (two-sided t-test p-value of 0.14). This does not necessarily indicate that there is no effect of belief elicitation on game play, since there are two remarks to make. First, the aggregate measure of average NE play is calculated across all games. But the hypotheses outlined above predict a stronger effect for games that are dominance solvable. While many of the 20 games are dominance solvable, not all are. Second, while the difference is not significant there are clear indications for a difference in the mean number of chosen NE actions depending on the treatment. Therefore, despite the difference being not significant, the means of all games could hide some effect of a subgroup of games. This is especially likely if one considers the relatively low power of the t-test in this setting.¹⁶

If only games that are dominance solvable are analyzed the picture changes and the effect size increases. While the percentage of equilibrium actions played increases in both treatments (to 61 percent for the baseline and to over 68 percent for the belief treatment), the difference between the two treatments is also much larger. A two-tailed t-test shows that this difference is significant at the one percent level (p-value 0.008). The result of the t-test confirms a significant difference between the treatments predicted by the two hypotheses described above. Belief elicitation increases the equilibrium play in games that are dominance solvable.

The effect becomes more apparent if one considers the histogram of chosen actions in accordance with a NE. Figure 1 states for each individual how often s/he played an action in accordance with a NE for all 13 dominance solvable games. The histogram is split for the two treatments. Since there are 13 decisions to make the maximal number of actions that can be in accordance with NE is 13. One observes that the right histogram, which shows the distribution for the belief treatment, seems to be “shifted” to the right compare to the left one. This is the effect on the mean of the distribution, which is detected by the t-test.¹⁷

The thinking harder hypothesis implies that a subject invests more effort thinking about the decision situation and therefore better understands the situation. Given this deeper understanding, it is expected that the subjects decide more like predicted by game theory. A consequence of this hypothesis for dominance solvable games is that a harder thinking subject is more likely to detect if the other player has a dominant or dominated action and reacts based on this knowledge. Therefore, if the thinking harder

¹⁵Of course the percentage highly depends on the games played. Rey-Biel (2009) for example finds higher rates of equilibrium play in arguable “easier” games, while Costa-Gomes and Weizsäcker (2008) find lower rates.

¹⁶The effect size in the data is $d \approx 0.26$. This Cohen’s d (Cohen, 1988) is defined as the difference between two means divided by the standard deviation of the given data. This is a fairly frequently used measure, especially in the psychological literature. With the sample size of 122 subjects, the power of the t-test is only around 25 percent (using GPower 3.1.4 (Faul et al., 2007)).

¹⁷Further restricting the set only to games which are strictly dominance solvable, meaning only including games that are dominance solvable by eliminating actions that are strictly dominated, yields similar results (p-value 0.0075). Dropping the observations from the crashed belief session yields the same qualitative results.




Figure 1: Histogram of NE actions for both treatments

hypothesis is true, I expect the subjects in the belief treatment to play actions that are iterative dominated, meaning actions that are only dominated if one eliminates all dominated actions of the other player from consideration, less often.

Surprisingly this effect cannot be found. Similar to the measure “NE actions chosen” above, I sum up the number of games in which a subject plays an action that is not iteratively dominated and compare these two aggregates for each treatment. There is no difference between the treatments.¹⁸ Therefore the chosen actions do not indicate that subjects are more likely to detect a dominated action of the other player in the belief treatment and respond accordingly.

The result is different if the players’ own dominated actions are considered. Playing a dominated action can be treated as committing an error, since a rational subject should never play a dominated action if he or she maximizes his / her own expected payoff. In the baseline treatment subjects play a dominated action on average in two games (mean 2.0). This is significantly larger than in the belief treatment, which has a mean of about 1.4 (two-sided t-test p-value 0.04). Hence it appears that belief elicitation leads to less play of own dominated actions. This can be taken as evidence for the “thinking harder” hypothesis. While on the one hand it is surprising that belief elicitation has an influence on finding own dominated actions, it is in line with a higher effort invested to think about the situation. The other explanation arguing that the effect on the chosen actions is due to the formation of beliefs on the other hand cannot explain this finding. Whether a subject has a dominated action does not depend on the chosen action of the other player, since an action that is dominated is never a best reply to any action of the other player independent of the possible payoff for the other player. Therefore it is an indication that supports the “thinking harder” hypothesis. My results hint that belief elicitation has a subject not only “think harder” about their likely chosen action of the other player, but also about their own best action. The “thinking harder” therefore is taking place on a lower level as expected, since it does not necessarily lead to more

¹⁸In the baseline treatment subjects play an iteratively dominated action in about 4.4 games, while in the belief treatment subjects choose such an action in about 4.2 games.

equilibrium behavior in all games, but to a less frequent choice of own dominated actions. This is a very interesting aspect of belief elicitation and sheds some light into the reasoning of subjects. Such a behavior is consistent with subjects not systematically determining their best response to all possible actions of the other player when they are not asked to state beliefs, but instead choose “where they would like to go” in the presented payoff matrix. When they are asked for beliefs they consider each possible action of the other player separately and therefore determine their best response for all possible actions of the other player. This results in less dominated actions to be chosen in the belief treatment.

4.2 Stated beliefs

Since the design of the experiment is such that subjects in both treatments state their beliefs, it is possible to compare the stated beliefs between the treatments. The first test concerns the question, whether subjects really do what they were asked to do and state their true beliefs. The major concern to address is that the used elicitation procedure with a quadratic scoring rule is only incentive compatible for risk-neutral subjects. Therefore the stated beliefs might be biased if many subjects do not state their true beliefs, but rather uniform belief distributions in order to secure themselves a sure payoff for the elicitation. If this is the case there should be many subjects who state relatively uniform belief patterns.

Analyzing the stated beliefs yields results which are not confirming these concerns. While some subjects seem to state uniform beliefs, the vast majority does not report beliefs that are even close to a uniform distribution. Figure 2 shows how often a subject stated beliefs that are not uniform. Uniform beliefs in these histogram are defined as stating no single belief above 51 percent if the other player has two actions to choose from and not stating a single belief above 34 percent if the other player has three actions to choose from. The histogram shows that one subject in the baseline treatment always stated uniform beliefs and a second does this in 17 out of 20 games. For those subjects the belief statements are clearly not the true beliefs. Either the subjects did not have any clue what the other player would be doing or they use the property of the quadratic scoring rule that they can get some payoff for sure by answering uniformly. While the last possibility seems to be the reason for a few subjects (these also answered the belief questions really quick), the most uniform belief statements should be from subjects who are simply unsure. Therefore the figure also shows that the problem of subjects only stating uniform beliefs is relatively limited. The vast majority of subjects does not report uniform beliefs.¹⁹ To analyze the individual stated beliefs systematically I categorized them into different categories. If a subject stated a belief for a game in which the opposing player has three actions to choose from I classify his or her belief into one of four possible categories. Category one contains all individual beliefs statements with at least a weight of 50% on the first action. In the second category are all individual beliefs of 50% or more on the second action of this game and similarly the third category contains all beliefs with a weight of at least 50% on the third action. If an individual stated beliefs that do not place a weight of at least 50% on an action his or her beliefs are categorized in the fourth category. For games in which the other

¹⁹Just to give some illustration for the distribution of individual beliefs there are two additional histograms in the appendix showing the number of stated beliefs for a single action that are higher than 60 respectively 70 percent. Pooling the numbers over both treatments, since the distributions appear to be very similar, shows that on average a subject puts a weight of 60 percent or more on a single action in about 9.7 games and a belief above 70 percent in more than seven games. These results indicate that most subjects do not state uniform beliefs.

player only has two available actions the cutoff is changed to 70% and there are only three instead of four categories. This classification allows me to explicitly test whether the stated beliefs are different from randomly chosen beliefs. As in the section above I first test this for each game separately and test for a difference in beliefs with respect to player roles. FEPT (5% significance level) yield only two rejections of equivalence in 40 comparisons.²⁰

While this confirms that there is no difference within treatment according to the player type, the question whether the stated beliefs are different from random beliefs remains. Using the same categorization as above and performing χ^2 goodness of fit tests yields for both treatments separately results in a rejection of random play for 19 games in both treatments. Only for game #2, in both treatments, the stated beliefs are indistinguishable from random guesses. Therefore I am confident to conclude that the stated beliefs are not different for the two player roles and are not the result of random answers.

The next question is whether or not the beliefs stated by the subjects in the two treatments are similar or not. There is some reason to suspect that the beliefs stated by the baseline subjects could be systematically different than those stated in the belief treatment. The reason for this is that if (very sophisticated) subjects in the baseline treatment expect an effect on game play due to the belief elicitation, they might state different beliefs. Given that the beliefs in the baseline treatment are elicited after all decisions have been made, a subject maximizes his / her expected payoff from the belief statement if s/he states beliefs that are as close as possible with the chosen distribution of actions. Therefore these subjects might state beliefs that take an effect of the elicitation on the chosen actions into account. While I consider this highly unlikely a difference in the stated beliefs between the two treatments could indicate such a behavior. Using the categorized beliefs and pooling the data across player roles within each treatment allows me to test for a treatment effect of the belief elicitation on the stated beliefs. FEPT conducted on each game separately yield no significant difference on the five percent level for any game.

These observations mitigate the concern that subjects in the baseline treatment respond differently to the question about their beliefs than subjects in the belief treatment. Therefore whatever drives the effect described in section 4.1 above, it seems to be triggered by the explicit belief elicitation. Therefore the stated beliefs are not different between the two treatments. If the belief elicitation would result in the described differences in chosen actions, but also to different stated beliefs in the two treatments one could argue that the belief elicitation in the baseline treatment are rationalizations of the own behavior in the part before. The fact that there is no difference in the stated beliefs is in line with both outlined hypotheses of the effect of belief elicitation on game play. Either the belief elicitation leads subjects to form beliefs or due to the belief elicitation subjects think more about the game. In both cases the fact that the beliefs are so similar in both treatments indicates that the effect is also working if the belief elicitation is done separately from the action choice.

4.3 Best response of actions to stated beliefs

In order to test the assumption that play is belief based I check whether the chosen actions are best responses to the stated beliefs. To test this I assume that subjects are maximizing their expected payoff.

²⁰Both rejections are in the belief treatment in games 8 and 14.




Figure 2: Histogram of not uniformly stated beliefs per subject

This assumption implies that subjects only care about their own monetary payoff and excludes any other motives for play like social preferences. I define a best response to be the action that maximizes the expected payoff of a subject given his or her stated beliefs.

Calculating the percentage of best responses over both treatments for all subjects yields a best response rate of 55.3 percent. This number is in line with the results of Costa-Gomes and Weizsäcker (2008) who also report best response rates around 50 percent and it is slightly lower than the results of Rey-Biel (2009). He finds best response rates around 67 percent. Given that his games are arguably easier than the ones used in my experiment the difference seems reasonable.²¹

Given that in the section above it was shown that the action choice is influenced by the belief elicitation, it is likely that the number of best responses is also influenced by the treatment. If belief elicitation leads to the formation of beliefs or to harder thinking, it is less likely that in the baseline treatment the chosen actions are best responses to the stated beliefs. When “thinking harder” is triggered by the explicit elicitation of beliefs and this process leads to a refinement of the beliefs and a likely consequence is a change in the best response. Once the belief changes also my answer and therefore my action choice might change. Given that the beliefs in the baseline treatment are not stated until all action choices are made and that the probability of NE play increases in the belief treatment, a difference in the number of best responses could arise. To test for such an effect I calculated the number of best responses

²¹I consider his games easier, since he has a higher share of dominance solvable games and half of the subjects only play constant sum games.




Figure 3: Histogram of best responses per subject to stated beliefs

per subject separately for both treatments. As the histogram in figure 3 indicates there is a difference between the two treatments.

On average subjects in the baseline treatment best respond to their own stated beliefs 9.9 times, while subjects in the belief treatment do so 12.2 times (out of 20 possible). The result of a t-test reveals that this difference in means is significant on the one percent level. A Kolmogorov-Smirnov test also rejects the equivalence of the distributions for the two treatments at the 1% level. Similar to the effect seen in figure 1 the t-test confirms the shift seen in the graph. The shift of the best response distribution seems to affect all subjects in the same way. The effect does affect subjects with a low and subjects with a high number of equilibrium actions in the same way.

It looks like subjects “think harder” as soon as they are asked to state their beliefs. When they realize that they did not notice an aspect of the game before, when they were choosing their action, e.g an own dominated action, they seem to believe that they are the only person who has overlooked this aspect. Therefore they state beliefs that are in line with their modified beliefs and not the beliefs they might have had when they actually made their decision. These results show that the belief elicitation triggers a modification of beliefs. The observation that the increased number of best responses in the belief treatment mainly affects subjects with a low level of equilibrium play shows that the effect of the belief elicitation seems to be at work in both treatments. While in the belief treatment it also leads to a higher level of equilibrium play, in the baseline treatment it results in modified beliefs. Given that I have shown above that there is no noticeable difference between the belief statements in both treatments, the

difference in the number of best responses is most likely explained by the fact, that the belief elicitation changes the beliefs of the subjects. In the belief treatment it also leads to an effect on the chosen actions, but this is impossible in the baseline treatment. The difference between the treatments in the number of best responses is a logical consequence. Since the belief elicitation leads to a modification of the beliefs it also influences the chosen actions (which are often best responses to these beliefs). Because subjects in the baseline treatment cannot react to their changed beliefs by choosing a different action. Consequently the number of best responses is lower in the baseline treatment. While I believe that this modification comes from an increased effort spend or more attention paid to the decision situation, it is clear that this reasoning is unlikely how subjects think about the process. It is not an explicit change of beliefs, rather the modification of beliefs is an unconsciously process, which yields a (slightly) different set of beliefs.

4.4 Level-k analysis

Since many experiments have shown that Nash equilibrium does not always predict how subjects decide various alternative models to the Nash equilibrium have be applied to explain the choices made by subjects. A very prominent one has been the level-k model of cognitive reasoning developed by Stahl and Wilson (1994). They develop a model with different player types that differ in their strategic thinking ability. Subjects with a low level of sophistication are assumed to play all of their available actions with equal probability. The *L1* model therefore predicts that a player plays all available actions with equal probability. A player who plays a *L2* strategy best responds to the predicted choice of a *L1* player and a *L3* player best responds to the play of a *L2* player. Higher levels of play, like *L4*, *L5* etc., are defined accordingly. Since previous studies have found that hardly any subject seems to play a strategy that is in accordance with a level of 4 or higher I will only consider the strategies *L1* to *L3*. The other models of game play which I consider are *D1* and *D2* play, *Maxmin* and *Efficiency*. The *D1* model of game play predicts that a player plays all of the available undominated actions with equal probability. *D2* players are assumed to best respond to the predicted choice of a *D1* player. A player that plays a *Maxmin* strategy always chooses the action with the highest minimum payoff, while an *Efficiency* (*Ef*) player chooses the action that maximizes the sum of both players payoff.

Table 2: Percentage of actions matched by models' predictions

Treatment	Games	NE	L1	L2	L3	D1	D2	Ef	Maxmin
Baseline	not DS	0.43	0.57	0.48	0.43	0.53	0.41	0.47	0.51
	DS	0.61	0.60	0.59	0.50	0.56	0.52	0.41	0.58
	All	0.55	0.59	0.55	0.47	0.55	0.48	0.43	0.56
Belief	not DS	0.39	0.58	0.50	0.40	0.55	0.44	0.50	0.57
	DS	0.68	0.60	0.65	0.56	0.61	0.59	0.33	0.62
	All	0.58	0.59	0.60	0.51	0.59	0.54	0.39	0.60
Average	All	0.56	0.59	0.58	0.49	0.57	0.51	0.41	0.58

Table 2 shows how many of the chosen actions are in accordance with the predictions of the various models. The table lists the percentages split for the two treatments baseline and belief, and further contains the results only for dominant solvable games (DS) and games that are not dominance solvable

(not DS) separately. The last row states the percentages pooled for all subjects for both treatments and all games. The table illustrates again that more actions in accordance with Nash equilibrium have been chosen in dominant solvable games. For all splits listed in the table the percentage of NE actions is higher for DS games compared to not dominance solvable games and, as discussed in section 4.1 above, the effect is larger in the belief treatment. Considering all choices in all treatments, *L1*, *L2* and *Maxmin* are slightly better than the *NE* predictions, but the differences are not large. The worst performing model is the *Efficiency* model, which does extremely poor for dominance solvable games especially in the belief treatment. It is important to keep in mind that for these two and three action games the probability for predicting the right choice is above 33% even for purely random predictions.²²

While between the treatments the number of choices that are correctly predicted by the *L1* model are roughly similar (always within one percentage point), the difference for the *L2* and *D2* predictions is larger. More subjects in the belief treatment make choices in accordance with the *L2* or the *D2* model compared to the subjects in the baseline treatment. Interestingly the difference for the two models comes from different games. As for the difference in NE actions chosen the difference between the treatments mainly stems from games which are dominance solvable. Subjects in the belief treatment make more choices in line with predictions from models with a higher levels of cognitive reasoning. This result nicely corresponds to the findings above that also imply that the belief elicitation increases the awareness of subjects and leads to more strategic sophistication.

The paragraphs above consider each choice of a subject separately. In the remainder of this section I analyze the individual behavior. To do this I classify subjects according to their individual choices in all 20 games jointly. A subject is said to be a “Nash equilibrium player” if NE predicts more chosen actions correctly than any other considered model. In the case that two or more models predict equally many choices correctly and no other model is better, the subject is classified as “not identified”. This classification leads to about a quarter of all subjects not being sorted into one of the eight categories. For the eight models considered here there is only one notable difference between the two treatments. When all games are considered the biggest difference between the two treatments is the portion of subjects that are classified as *Efficiency* types. Many more subjects are classified as *Efficiency* types in the baseline compared to the belief treatment. There is no difference in the percentage of subjects classified in any other category.²³ Based on the results reported in section 4 I also classify subjects only based on their choices in the 13 dominance solvable games. I do this, since the previous sections have illustrated that an effect of the belief elicitation seems to be present for exactly these dominance solvable games. If there exists a difference in the choices of subjects between the two treatments it is likely that this difference is most pronounced in the dominance solvable games and therefore I expect this difference to show up in the classification performed in this section. By restricting the classification to these games it becomes more likely that the simply sorting mechanism used here (the modal number of correct predictions for each subject) can identify a difference between the treatments.

²²If always three actions would be possible in each game and subjects would choose their actions randomly, a model should predict the correct choice in 33% of all cases. Therefore the 39% of the *Efficiency* model are really weak.

²³The difference for the *Efficiency* predictions is large with more than 15% classified as *Efficiency* types in the baseline, but less than 4% in the belief treatment. For details of the analysis with all games see the complete table 7, which is given in the appendix.

Table 3: Percentage of subjects classified by models' actions most taken - only dominant solvable games

Model	Clear Cases		Ties		Clear cases and ties	
	Baseline	Belief	Baseline	Belief	Baseline	Belief
NE	13.33	29.03	3.33	1.61	16.66	30.64
L1	13.33	14.52	5.00	1.61	18.33	16.13
L2	11.67	12.90	5.00	1.61	16.67	14.51
L3	0.00	0.00	0.83	0.00	0.83	0.00
D1	1.67	3.23	4.17	1.61	5.84	4.84
D2	0.00	0.00	0.00	0.81	0.00	0.81
Ef	15.00	4.84	2.50	0.81	17.50	5.65
Maxmin	15.00	14.52	2.50	1.61	17.50	16.13
not identified	30.00	20.97	6.67	11.29	6.67	11.29

Notes: The second and third column only states unique predictions. The ties, which are reported in the fourth and fifth column, are ties between exactly two models. In case of a tie the subjects is classified as 50% for each model.

Table 3 states the percentage of subjects that can be classified according to the different models if only the 13 dominance solvable games are used for the classification. *Clear cases* are subjects for which one model predicts more choices correctly than any other model of game play considered here. All other subjects are listed as *not identified*. *Ties* are only stated for subjects for which exactly two different models predict the most chosen actions correctly. In this case a subjects is split between the two models which predict equally well. Therefore, if three or more models predict equally many choices of a subject correctly, and no other model predicts more choices correctly, the subject is still classified as *not identified* in columns four and five. The last two columns lists the total percentage of subjects for each model split for the two treatments, including clear cases and ties between two models.

The table shows that more subjects are classified as “Nash” players in the belief treatment compared to the baseline treatment. Almost twice as many subjects are classified as “Nash” players in the belief treatment. The 30% of subjects classified as “Nash” players is also the largest group within the belief treatment. Compared to all other seven models about twice as many subjects are in this group. The difference between the treatments is to be expected, given that we have seen before that in the belief treatment many more choices are in line with the prediction of Nash equilibrium. The only other notable difference is for the *Efficiency* model. While 17.5% of subjects in the baseline treatment are classified as *Efficiency* types less than 6% are in this category for the belief treatment. The percentages in the other categories are rather similar between the two treatments. The number of subjects that cannot be classified is much higher for the belief treatment once ties between two models are considered. For some reason, which is not obvious to me, there are many ties between exactly two models in the baseline treatment, but many more ties between three or more models in the belief treatment.

This analysis yields two major findings. The first is a confirmation of the findings in other sections. Belief elicitation affects game play for dominance solvable games and leads to more equilibrium play. Consequently more subjects are classified as *Nash equilibrium players* in the belief than in the baseline treatment. Second, at least for the games considered here, no other model of cognitive reasoning outperforms the Nash predictions. Also there is no difference between the treatments, besides the difference in “Nash” types. Therefore this analysis cannot support the claim, that belief elicitation mainly works

on subjects with a relatively low level on strategic thinking. If this would be the case, there should be a difference in the proportion of subjects classified into “lower levels” of strategic thinking, such as *L1* or *D1*, between the treatments. As can be seen in table 3 such a difference is not found.

When discussing the difference considering all choice as well as the classification based on the dominance solvable games, it is important to keep in mind that the experiment was not mainly designed to test for differences between the models listed here. This means that different than in studies that focus on cognitive reasoning models the games are not selected to produce large differences between the predictions of the models. Out of the 20 games, the NE and the *L1* model therefore predict different actions to be chosen in nine games. Naturally the number of games for which the predictions are different to NE are even lower for models of “higher cognitive reasoning” like the *L2* or the *D2* model. There are four games in which the predictions for the chosen action are different between the *L2* and the NE model and three games for the predictions made by the *D2* and the NE model. Therefore it is remarkable to see that there are still the reported differences between the models success in predicting choices. If the games would be selected with the main purpose of differentiating between these model, I would expect even larger differences.

4.5 The influence of various game properties on chosen actions

The analysis in the previous sections has focused on the general effect of belief elicitation on game play using aggregate measures over all games for an individual subject. The only property of the games considered was whether or not a game is dominance solvable. In this section I make use of the variety of other properties that differ between the games. In order to exploit the various game properties no aggregate measures for each individual are used, but rather each decision (action choice) of an individual is used separately. The variable of interest is whether the chosen action is part of a pure strategy NE or not. Since each subject makes 20 decisions there are in total 2440 observations (122 subjects) available for the analysis. The main properties in which the used games differ, besides whether a game is dominance solvable, are whether a game is symmetric or not, whether an symmetric alternative to the NE exists and whether the NE is Pareto dominated by another outcome. Given the mixed results found in the literature about a possible effect of belief elicitation on game play it is especially interesting to test which game properties influence the effect. Similar to the research about the likelihood of Nash equilibrium play, in which various properties of the game and the decision environment have been shown to influence the predictive power of Nash equilibrium, it is likely that the effect of belief elicitation also depends on some game properties. Before it was shown that “easier” games, meaning dominance solvable games, are affected, while over all the effect was not significant. Therefore in this section the influence of other game properties is investigated and the previous result are tested with a different measure.

To detect any effect of these game properties on the likelihood that an action in accordance with a NE is played I run a simple linear probability model on dummies for these four properties: dominance solvability (DS), symmetry (symmetric), whether a symmetric alternative to the NE exists (symmetric_alt) and whether the NE is Pareto dominated by another outcome (NE_Pareto_dominated). Also a treatment dummy for decisions made in the belief treatment is used. This dummy variable is one for all decisions that are made in the treatment in which subjects simultaneously state their action choice and

their beliefs and zero otherwise. Table 4 states the results of this OLS model.

Table 4: Linear probability estimation

Variable	Coefficient	(Std. Err.)
treatment	0.034	(0.023)
DS	0.468***	(0.025)
symmetric	0.188***	(0.035)
NE_Pareto_dominated	-0.441***	(0.025)
symmetric_alt	-0.124***	(0.025)
Intercept	0.415***	(0.029)
<hr/>		
N	2440	
R ²	0.172	
F (5,121)	109.746	

Significance levels : * : 10% ** : 5% *** : 1%

Standard errors are cluster on the subject level.

The results are very much in line with the analysis of the aggregate data above. When all games are considered there is no treatment effect on the likelihood of equilibrium play. Accordingly the coefficient on the treatment dummy is not significantly different from zero. This is not the case for the game properties. As one might expect after seeing the results above, the probability of NE play is significantly increased in games which are dominance solvable or symmetric. The coefficient for both dummy variables (DS and symmetric) are significantly positive at the one percent level (p-values < 0.01), with the coefficient for the dominance solvability being more than two times as large as the coefficient for symmetric games. Also the two other coefficients are significant and point in the expected direction. NE play is less likely if there exists an alternative symmetric outcome to the NE or if the NE is Pareto dominated. Overall, this regression replicates the finding from the treatment analysis above and further confirms that various game properties have an influence on the likelihood of NE play. Since in this analysis all decisions of an individual are considered separately it is very likely that the residuals are not homogeneous, as the errors of the same individual are most likely correlated. Therefore the shown regression, and all that follow below, include standard errors, which are clustered at the subject level.²⁴

In order to control whether there exists a treatment effect for a subpopulation of games, the next regression additionally includes interaction effect between the treatment dummy and the variable for dominance solvability (DS). This regression repeats the analysis made with the aggregate measure above. The dummy is one if a game is dominance solvable and the decision is made in a session in which beliefs and actions are stated at the same time. As table 5 shows the inclusion of this interaction terms does not change the sign or significance of the coefficients used in both regressions. The interaction effect between the treatment dummy and the dominance solvability dummy is significantly positive at the one percent level. Therefore the likelihood that a subject chooses an action in accordance with NE is even larger in games that are dominance solvable when beliefs are elicited together with the action choice.

²⁴The results are unchanged if a probit or logit model instead of the linear regression is used. In the appendix the results of a probit estimation with marginal effects evaluated at the mean are shown. Also the probit regression has standard errors clustered at the subject level.

While the size of the interaction effect is modest, compared to the effect of dominance solvability alone, it is important to note that this effect basically “is on top”. The analysis of the individual games therefore confirms the finding from before. There exists a treatment effect of the belief elicitation on the likelihood of NE play.²⁵

Table 5: Linear estimation with interaction effect

Variable	Coefficient	(Std. Err.)
treatment	-0.039	(0.037)
DS	0.410***	(0.034)
symmetric	0.188***	(0.035)
NE_Pareto_dominated	-0.441***	(0.025)
symmetric_alt	-0.124***	(0.025)
int_DS_treatment	0.112***	(0.043)
Intercept	0.452***	(0.033)
<hr/>		
N	2440	
R ²	0.175	
F (6,121)	95.562	

Significance levels : * : 10% ** : 5% *** : 1%

All standard errors are cluster on the subject level.

The effect of the interaction between the treatment and dominance solvable games is very robust. Restricting the analysis only to games that are dominance solvable yields qualitatively similar results.²⁶ The treatment coefficient remains positive and significant. Also the other coefficients of the game properties also remain significant. The only change concerns the coefficient on the symmetric alternative. When only dominance solvable games are considered the coefficient becomes significantly positive. While this is clearly counterintuitive the most likely explanation is the fact that only one pair of games is dominance solvable and has a symmetric alternative to the NE. Therefore the coefficient basically measures the difference of the likelihood of NE play in these two games relatively to all other games. It seems to be the case the NE is played more often in these games compared to the average in all the other games. Certainly this is not enough information to conclude a general pattern. The analysis is also robust to the used regression model. Even though all results mentioned so far are based on a linear probability model they remain unchanged if a binary choice model, e.g. a probit or logit regression, is used instead. All coefficients are evaluated at their respective means. As for the linear probability model all regression are done with standard errors clustered at the subject level.

To test whether the effect of belief elicitation also depends on other game properties table 6 contains a linear regression with four interaction effects. All game properties are interacted with the treatment dummy. This means that the “baseline category” are games that are not dominance solvable, have no symmetric alternative to the NE, have a NE that is not Pareto dominated by any other outcome and are not symmetric. Additionally all regressors of the estimations before are included in the estimation. The

²⁵ Again see the appendix for a probit estimation which yields similar results.

²⁶ All regression results for the subgroup of dominance solvable games are reported in table 10 which is provided in the appendix.

Table 6: Linear probability estimation with interaction effects

Variable	Coefficient	(Std. Err.)
treatment	0.033	(0.049)
DS	0.422***	(0.036)
symmetric	0.275***	(0.048)
NE_Pareto_dominated	-0.428***	(0.035)
symmetric_alt	-0.068*	(0.039)
int_DS_treatment	0.089*	(0.049)
int_sym_treatment	-0.171**	(0.068)
int_sym_alt_treatment	-0.111**	(0.049)
int_Pareto_treatment	-0.026	(0.050)
Intercept	0.416***	(0.038)
<hr/>		
N	2440	
R ²	0.176	
F _(9,121)	68.549	

Significance levels : * : 10% ** : 5% *** : 1%

All standard errors are cluster on the subject level.

results reveal that the effect of the dominance solvability, the symmetric and whether a Pareto improvement to the NE outcome exists have the same effect as before and all remain highly significant. Only the effect on the presence of a symmetric alternative is weakened. It is now no longer significant at the five percent level. Regarding the interaction terms the biggest surprise is the sign change of the dummy for the interaction of treatment and the dominance solvability. While the coefficient estimate is still positive it is not significantly different from zero on the five percent level. Instead the two interactions between the treatment dummy and the symmetric of the game and the existence of a symmetric alternative are both negative and significantly different from zero. While the relative small magnitudes of all this interaction effects hints that the differences are not too large this result comes as a surprise. Especially that the likelihood of NE play is decreased, according to the interaction effect, if the game is symmetric and the choice is made in the belief treatment is unexpected. Nevertheless this analysis shows that not all games are affected similarly by the elicitation.

Overall the analysis of the individual game properties results in two main findings. First, the treatment effect, which was before shown for aggregated measures on the subject level, can also be identified using all decisions separately. The belief elicitation does increase the likelihood of NE actions in dominance solvable games. Second, the results of the game properties are in line with other results in the literature. Without specifying a complete model of game complexity it seems as if NE play is more likely in games that are relatively “easy”, meaning in games that are dominance solvable and symmetric. The fact that the treatment effect can only be found in the subgroup of dominance solvable games hints, that there is an effect of belief elicitation, but that it is too small to show up in more complicated games. These findings are in line with the “thinking harder” hypothesis and therefore this section provides further evidence that shows that subjects seem to invest more effort to think about a decision situation when they are asked to state their beliefs explicitly. This yields an increased play of NE actions in relatively easy games.

5 Discussion and conclusion

In this paper I have shown that belief elicitation can have a significant effect on the chosen actions of individuals. Asking for the beliefs of subjects together with their action choice increases equilibrium play in dominance solvable games and reduces the number of dominated actions played. My analysis shows that belief elicitation does not affect the choices in all games equally. The properties of the game play an important role for the presence and the size of the effect. While it is surprising that the effect is driven by the increased recognition of own dominated actions and the subsequent avoidance of these actions this observation is consistent with a deeper understanding of the game due to the belief elicitation. Therefore my experiment reveals that there exists a treatment effect of belief elicitation and that the elicitation facilitates the understanding of dominated actions. This result supports the so-called “thinking harder” hypothesis, which claims that belief elicitation raises awareness of subjects for the situation or has subjects increase their effort to understand the game. Therefore they understand the decision situation better and this results in choices which are closer to the game theoretic predictions.

When analyzing the stated beliefs there is no difference between the treatments, but a huge difference between the number of best responses to the own stated beliefs. The fact that in the belief treatment subjects respond more frequently optimally to their own states beliefs is not surprising, but fits into the story of “thinking harder”. Investing more effort in understanding the game and especially the likely behavior of the other subjects does not only lead to less dominated actions being played, but also to more best responses. The setup of the experiment leads subjects to revise their beliefs in the baseline treatment, when they are explicitly asked to state these beliefs. Since they cannot change their action choices afterwards the chosen actions are less often best responses to the stated beliefs compared to the belief treatment.

These results suggests that belief elicitation can be used as an experimental tool. Even if a researcher is not particular interested in the stated beliefs themselves, the effect that subjects are “positively” affected by the elicitation could make the elicitation of beliefs worthwhile. In general, the experiment is not able to definitely distinguish between the “thinking harder“ hypothesis and the claim that belief elicitation leads to the formation of beliefs, since some results can also be explained by the second hypothesis. But the effect of a decreased play of own dominated actions and the difference in the likelihood of best response between the two treatments clearly point towards to the “thinking harder” hypothesis. Therefore the experiment provides clear evidence for an effect of belief elicitation on game play and additionally provides several indications why this is the case.

An aspect which should be analyzed further is the role of the difficulty or complexity of the games. The analysis in the previous sections has shown that game properties have an influence on the likelihood of equilibrium play and on the effect of belief elicitation on game play. Belief elicitation seems to enhance the understanding of the task for the subjects, but only if the game is not “too complicated”. Therefore, belief elicitation matters only for games that are not too difficult. Given the mixed results of other papers on the effect of belief elicitation the “complexity” of the games might be a logical explanation. Assuming the “thinking harder” hypothesis is true, it is very likely that belief elicitation does not play a major role in relatively easy games (subjects understand the game anyhow) nor in very complicated games (in which the effect is not strong enough). The fact that the treatment effect was only

significant for dominance solvable games clearly points into this direction. In further studies it would be interesting to test whether other game properties, besides dominance solvability, also play a role and if the effect can be extended to other situations besides normal-form games.

References

- Bernheim, B. Douglas**, “Rationalizable Strategic Behavior,” *Econometrica*, 1984, 52 (4), pp. 1007–1028.
- Blanco, Mariana, Dirk Engelmann, Alexander Koch, and Hans-Theo Normann**, “Belief elicitation in experiments: is there a hedging problem?,” *Experimental Economics*, 2010, 13, 412–438.
- Camerer, Colin F.**, *Behavioral game theory: Experiments in strategic interaction*, Princeton University Press, 2003.
- Cohen, Jacob**, *Statistical power analysis for the behavioral sciences*, Lawrence Erlbaum, 1988.
- Costa-Gomes, Miguel A. and Georg Weizsäcker**, “Stated Beliefs and Play in Normal-Form Games,” *The Review of Economic Studies*, 2008, 75 (3), pp. 729–762.
- Costa-Gomes, Miguel, Vincent P. Crawford, and Bruno Broseta**, “Cognition and Behavior in Normal-Form Games: An Experimental Study,” *Econometrica*, 2001, 69 (5), pp. 1193–1235.
- Croson, Rachel T.A.**, “The Disjunction Effect and Reason-Based Choice in Games,” *Organizational Behavior and Human Decision Processes*, 1999, 80 (2), 118 – 133.
- , “Thinking like a game theorist: factors affecting the frequency of equilibrium play,” *Journal of Economic Behavior & Organization*, 2000, 41 (3), 299 – 314.
- Dufwenberg, Martin and Uri Gneezy**, “Measuring Beliefs in an Experimental Lost Wallet Game,” *Games and Economic Behavior*, 2000, 30 (2), 163 – 182.
- Faul, Franz, Edgar Erdfelder, Albert-Georg Lang, and Axel Buchner**, “G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences,” *Behavior Research Methods*, 2007, 39, 175–191.
- Fischbacher, Urs**, “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, 2007, 10 (2), 171–178.
- and **Simon Gächter**, “Social Preferences, Beliefs, and the Dynamics of Free Riding in Public Goods Experiments,” *American Economic Review*, 2010, 100 (1), 541–56.
- , —, and **Ernst Fehr**, “Are people conditionally cooperative? Evidence from a public goods experiment,” *Economics Letters*, 2001, 71 (3), 397 – 404.
- Gächter, Simon and Elke Renner**, “The effects of (incentivized) belief elicitation in public goods experiments,” *Experimental Economics*, 2010, 13, 364–377.
- Greiner, Ben**, “An Online Recruitment System for Economic Experiments,” in “Forschung und wissenschaftliches Rechnen,” Kurt Kremer, Volker Macho, 2004, pp. 79–93.
- Haruvy, Ernan, Yaron Lahav, and Charles N. Noussair**, “Traders’ Expectations in Asset Markets: Experimental Evidence,” *The American Economic Review*, 2007, 97 (5), pp. 1901–1920.

- Huck, Steffen and Georg Weizsäcker**, “Do players correctly estimate what others do?: Evidence of conservatism in beliefs,” *Journal of Economic Behavior & Organization*, 2002, 47 (1), 71 – 85.
- Kagel, John H. and A.E. Roth**, *The handbook of experimental economics*, Princeton university press, 1995.
- McKelvey, Richard D. and Talbot Page**, “Public and Private Information: An Experimental Study of Information Pooling,” *Econometrica*, 1990, 58 (6), pp. 1321–1339.
- Nyarko, Yaw and Andrew Schotter**, “An Experimental Study of Belief Learning Using Elicited Beliefs,” *Econometrica*, 2002, 70 (3), pp. 971–1005.
- Offerman, Theo, Joep Sonnemans, and Arthur Schram**, “Value Orientations, Expectations and Voluntary Contributions in Public Goods,” *The Economic Journal*, 1996, 106 (437), pp. 817–845.
- Pearce, David G.**, “Rationalizable Strategic Behavior and the Problem of Perfection,” *Econometrica*, 1984, 52 (4), pp. 1029–1050.
- Rey-Biel, Pedro**, “Equilibrium play and best response to (stated) beliefs in normal form games,” *Games and Economic Behavior*, 2009, 65 (2), 572 – 585.
- Rutström, E. Elisabet and Nathaniel T. Wilcox**, “Stated beliefs versus inferred beliefs: A methodological inquiry and experimental test,” *Games and Economic Behavior*, 2009, 67 (2), 616 – 632.
- Schlag, Karl, James Tremewan, and Joel van der Weele**, “A Penny for Your Thoughts: A Survey of Incentives and Belief Elicitation,” *Manuscript, Frankfurt University*, 2013.
- Stahl, Dale O. and Paul W. Wilson**, “Experimental evidence on players’ models of other players,” *Journal of Economic Behavior & Organization*, 1994, 25 (3), 309 – 327.
- Weizsäcker, Georg**, “Ignoring the rationality of others: evidence from experimental normal-form games,” *Games and Economic Behavior*, 2003, 44 (1), 145 – 171.
- Wilcox, N.T. and N. Feltovich**, “Thinking like a game theorist: Comment,” *University of Houston Department of Economics working paper*, 2000.
- Zelmer, Jennifer**, “Linear Public Goods Experiments: A Meta-Analysis,” *Experimental Economics*, 2003, 6, 299–310.

Appendix

Games 1-4 are symmetric. All even games from game number six on are isomorphic transformations of the previous game. In those games the numbers are slightly changed and the order of actions is modified. In all treatments and sessions the order of games was identical and the following: 7, 11, 2, 8, 12, 19, 1, 6, 20, 5, 10, 13, 17, 9, 18, 14, 4, 16, 3 and 15. This sequence ensures that equivalent games are never played directly after each other. Pure Nash equilibria outcomes are underlined.

Game 1

	'T'	'M'	'B'
T	40, 40	20, 30	0, 20
M	30, 20	20, 20	100, 10
B	20, 0	10, 100	<u>100, 100</u>

Game 2

	'T'	'B'
T	75, 75	25, 85
B	85, 25	<u>30, 30</u>

Game 3

	'T'	'M'	'B'
T	10, 10	49, 58	37, 60
M	58, 49	22, 22	16, 36
B	60, 37	36, 16	<u>38, 38</u>

Game 4

	'T'	'M'	'B'
T	<u>26, 26</u>	22, 17	18, 24
M	17, 22	11, 11	33, 21
B	24, 18	21, 33	20, 20

Game 5

	L	R
T	<u>55, 79</u>	84, 52
B	31, 46	72, 93

Game 6

	L	R
T	48, 80	89, 68
B	<u>75, 51</u>	42, 27

Game #6's payoffs are obtained by subtracting four points from Game #5's payoffs.

Game 7

	L	C	R
T	74, 38	78, 71	46, 43
M	96, 12	10, 89	57, 25
B	15, 51	83, 18	<u>69, 62</u>

Game 8

	L	C	R
T	73, 80	20, 85	91, 12
M	45, 48	<u>64, 71</u>	27, 59
B	40, 76	53, 17	14, 98

Game #8's payoffs are obtained by adding two points from Game #7's payoffs.

Game 9

	L	C	R
T	31, 25	23, 12	<u>42, 33</u>
M	36, 36	30, 20	18, 27
B	48, 31	29, 42	23, 16

Game 10

	L	C	R
T	29, 20	18, 25	<u>35, 44</u>
M	22, 32	44, 31	14, 25
B	38, 38	33, 50	27, 33

Game #10's payoffs are obtained by adding four points from Game #9's payoffs.

Game 11

	L	C	R
T	52, 23	36, 27	16, 41
M	61, 31	31, 46	22, 28
B	80, 15	12, 40	34, 53

Game 12

	L	C	R
T	49, 30	24, 18	37, 12
M	11, 76	27, 57	19, 48
B	36, 8	42, 27	23, 32

Game #12's payoffs are obtained by subtracting four points from Game #11's payoffs.

Game 13

	L	C	R
T	43, 68	11, 31	23, 31
M	38, 38	44, 25	51, 38
B	31, 12	52, 29	45, 33

Game 14

	L	C	R
T	33, 13	27, 46	31, 54
M	33, 25	40, 53	35, 47
B	70, 45	40, 40	14, 33

Game #14's payoffs are obtained by adding two points from Game #13's payoffs.

Game 15

	L	R
T	74, 62	43, 40
M	25, 12	76, 93
B	59, 37	94, 16

Game 16

	L	C	R
T	10, 23	35, 57	60, 72
B	91, 74	14, 92	38, 41

Game #16's payoffs are obtained by subtracting two points from Game #15's payoffs.

Game 17

	L	C	R
T	67, 46	43, 31	61, 16
B	32, 86	52, 52	89, 62

Game 18

	L	R
T	42, 63	82, 28
M	12, 57	58, 85
B	27, 39	48, 48

Game #18's payoffs are obtained by subtracting four points from Game #17's payoffs.

Game 19

	L	C	R
T	32, 29	26, 18	39, 39
B	19, 28	45, 32	42, 19

Game 20

	L	R
T	30, 21	31, 34
M	21, 44	41, 41
B	34, 47	20, 28

Game #20's payoffs are obtained by adding two points from Game #19's payoffs.

Appendix II

The quadratic scoring rule was used for the payment of the belief statements. The exact formula is given below and for both treatments the following parameters were used: $\alpha = 80$ and $\beta = 0.004$. These points are converted to Euros at a rate of 20 points = 1 Euro. I_k is set to 100 and the belief statements (p_k) are between 0 and 100. The formula is identical for all games and treatments.

$$S_i(p) = \alpha - \beta \sum_{k=1}^n (I_k - p_k)^2$$

- I_k indicator function
- p_k - stated belief that the event k takes place
- $k = 1, 2, \dots, n$ all possible events
- $\alpha > 0$ - maximal payoff
- $\beta > 0$ - penalty term

Appendix III

Below are the English translations of the original German instructions. The original instructions are available from the author upon request. The instructions are for the belief treatment. The instructions for the baseline treatment are almost identical, but split in two parts (games and belief statements). The emphases are also in the original instructions. The first part (WELCOME) was given to the participants at the start, while the second part (Detailed instructions) was handed out after all participants were ready to start the actual experiment.

WELCOME!

Thank you for decision to take part in this experiment. By following the instructions you can, depending on your own and the decisions of others in the experiment, earn money. The entire experiment will last about **90 minutes** and you will receive your payment in cash directly after the experiment.

It is very important that you remain silent and do not talk with other participants during the entire experiment. In case you have any questions regarding the procedures of the experiment, please raise your hand and an experimenter will come to you to clarify the question directly with you. Thank you for your cooperation.

During the experiment you will be matched with other participants. We will always call the other participant “person B”. In each round you will be matched with a new person, this means you will never be matched with the same person in two successive rounds. Neither you nor person B will during or after the experiment know in which round you were matched with each other. Both of you will simultaneously see a decision situation. Your decision and the decision of person B together yield the result. This result determines how many points you and person B will receive and therefore how high your payoff is going to be.

We will now provide you the detailed description of the decision situation. If you have any questions, please raise your hand. We will come to your seat and answer open questions.

Detailed instructions for the decision situation

In this experiment we will show you multiple decision situations, an example can be seen in the table below. Please take a good look at the table and read the instructions below.

	Person B: §	Person B: #	Person B: +
You: &	40 10	40 40	10 10
You: %	50 20	40 40	0 20
You: =	20 30	20 50	40 10

In the real situations you will see a table like this on your computer screen (but with different numbers and possibly less cells). You will be asked to make a decision, which means that you chose between the given alternatives (here §, % and #). Person B decides simultaneously between their alternatives (here !, = and §). Therefore you always decide between the rows and person B always between the columns of the table. The cell that results from the combination of decisions states the results and therefore the number of points for you and person B. Your potential payoffs are given in each cell in the bottom left, while the potential payoff for person B is given in the upper right of each cell. As an example the following payoffs are possible:

- You chose § and person B !, **you** will get 10 and person B 40 points.
- You chose % and person B !, **you** will get 20 and person B 50 points.
- You chose # and person B =, **you** will get 50 and person B 20 points.

You will be encounter multiple of those decision situations during the experiment. At the end of the experiment three situations will be chosen at random and you will receive the sum of points, which you earned in these three selected situations. All points you earned in these three situations add up to your total points. **For each point you will receive 5 cents at the end of the experiment.** This means 20 points are exactly to 1 Euro.

Furthermore we will asked you about your expectation, which alternative person B will most likely choose. This means in each situation we will ask you to think about the behavior of person B. You can interpret the question as follows: If 100 persons would make this decision, how often would each of the given alternatives be chosen? Of course only one person will actually make the decision and not 100 different persons. But you can interpret the question about the behavior as how likely it is that a person B will choose each of the possible alternatives (here: !, = and §).

For example, if you are certain that person B will always select alternative ! in a decision situation and never alternatives = or §, then you would answer the question by stating the numbers 100, 0 and 0 for the fields of the alternatives !, = and §.

If you instead believe that person B will not always chose alternative !, but alternative = more often and action § most likely the most, then you can for example enter the numbers 20, 30 and 50 for the respective alternatives. If you assume that all alternatives are equally likely you can enter the numbers 33, 34 and 33 for the respective fields.

WARNING: The sum of all expectations always has to sum up to 100.

Besides the three situations in which the result of the decision will be paid, three additional situations will be randomly selected and you will receive the sum of points, which you have earned for your expectations in these three rounds. All points you earn in these decision situations are added to your total points.

For each decision situation only the result **OR** your expectations can be paid. This means, you will be paid for a total of six decision situations. In three situations for the result and in three different situations for your expectations. You will never receive points for both tasks in the same situation.

The points total which you receive for your expectations is calculated with a given formula. This formula is the same for each decision situation and is chosen in such a way that you maximize your total expected payoff if you state your true expectations. The table below provides four examples about the size of your payment.

The first table shows four examples of possible expectations about the chosen decision by person B. As an example row 1 states the following expectation: You believe that person B will always chose alternative ! and never the alternatives = or §.

	Person B: §	Person B: #	Person B: +
Example 1:	100	0	0
Example 2:	20	30	50
Example 3:	33	34	33
Payoff in example 1:	50	0	0
Payoff in example 2:	26	31	41
Payoff in example 3:	33	34	33

The next table states how large your payment would be, if person B actually choses alternative !, = or §. The first column states the payment you receive if person B choses alternative !, in the second column the payoff if s/he selects alternative = and the third column states your payoff if s/he choose action §.

If you state in your expectation as in example one and person B selects alternative !, then you receive 80 points. Does person B not choose this alternative, but = or §, then you receive 0 points.

If you state your expectations like in example 3 and person B choses alternative ! you get 53 points. Does person B choose alternative = you get 54 points and if alternative § is selected you receive 53 points.

The table states for the four examples all possible payoffs, depending on the decision of person B. **Each point is worth 5 cents**, independently if you received it through the decision or your expectation. You can maximally get 80 point and minimal 0 points for an expectation. In the example this means when person B selects alternative ! and you are paid for your expectation you would receive 80 points (= 4 Euro) in the first, 41 points (= 2.1 Euro) in the second and 53 points (= 2.7 Euros) in the third example. It maximizes your expected payoff, if you always state your true expectations.

Warning: The numbers given here are only examples. These examples do NOT provide any suggestions for the decision situations.

Please make sure that you understood the table and that you know how you receive points for your stated expectations. Raise your hand if anything is unclear.

Once you have the feeling that you understood everything and you now know how you are paid in these decision situations, please click with the mouse on the OK button. We will then ask you some understanding questions to a different decision situation. After you have answer these questions correctly you can start with the experiment.

Appendix IV

Table 7: Percentage of subjects classified by models' actions

Model	Clear Cases		Ties		Clear cases and ties	
	Baseline	Belief	Baseline	Belief	Baseline	Belief
NE	10.00	9.68	2.50	1.61	12.50	11.29
L1	16.67	19.35	5.00	0.81	21.67	20.16
L2	15.00	11.29	3.33	3.23	18.33	14.52
L3	0.00	0.00	0.00	0.00	0.00	0.00
D1	5.00	8.06	1.67	0.81	6.67	8.87
D2	0.00	3.23	0.83	0.00	0.83	3.23
Ef	13.33	3.23	2.50	0.00	15.83	3.23
Maxmin	16.67	20.97	2.50	3.23	19.17	24.20
not identified	23.33	24.19	5.00	14.51	5.00	14.51

Appendix V

This section contains additional estimation results for the analysis of section 4.5 on the influence of different game properties on the likelihood of equilibrium play.

Table 8: Probit estimation

Variable	Coefficient	(Std. Err.)
treatment	0.103	(0.068)
DS	1.328**	(0.086)
symmetric	0.497**	(0.105)
NE_Pareto_dominated	-1.255**	(0.087)
symmetric_alt	-0.342**	(0.077)
Intercept	-0.224**	(0.080)
<hr/>		
N	2440	
Log-likelihood	-1449.609	
$\chi^2_{(5)}$	397.669	

Significance levels : * : 10% ** : 5% *** : 1%

All standard errors are clustered on the subject level.

Table 9: Probit estimation with interaction effect

Variable	Coefficient	(Std. Err.)
treatment	-0.109	(0.102)
DS	1.163**	(0.106)
symmetric	0.497**	(0.106)
NE_Pareto_dominated	-1.261**	(0.088)
symmetric_alt	-0.344**	(0.078)
int_DS_treatment	0.337**	(0.124)
Intercept	-0.116	(0.092)
<hr/>		
N	2440	
Log-likelihood	-1445.084	
$\chi^2_{(6)}$	402.493	

Significance levels : * : 10% ** : 5% *** : 1%

All standard errors are clustered on the subject level.

Table 10: Linear probability estimation (only DS games)

Variable	Coefficient	(Std. Err.)
treatment	0.073**	(0.027)
symmetric	0.537**	(0.043)
NE_Pareto_dominated	-0.342**	(0.029)
symmetric_alt	0.135**	(0.035)
Intercept	0.733**	(0.034)
<hr/>		
N	1586	
R ²	0.192	
F (4,121)	104.193	

Significance levels : * : 10% ** : 5% *** : 1%

All standard errors are clustered on the subject level.

