

University of Heidelberg

Department of Economics



Discussion Paper Series | No. 629

Does mutual knowledge of preferences lead to more
equilibrium play? Experimental evidence

Christoph Brunner, T. Florian Kauffeldt,
and Hannes Rau

January 2017

Does mutual knowledge of preferences lead to more equilibrium play? Experimental evidence[†]

Christoph Brunner[‡] and T. Florian Kauffeldt[‡] and Hannes Rau[‡]

This version: January 20, 2017

Abstract

In many experiments, the Nash equilibrium concept seems not to predict well. One reason may be that players have non-selfish preferences over outcomes. As a consequence, even when they are told what the material payoffs of the game are, mutual knowledge of preferences may not be satisfied. We experimentally examine several 2×2 games and test whether revealing players' preferences leads to more equilibrium play. For that purpose, we elicit subjects' preferences over outcomes before the games are played. It turns out that subjects are significantly more likely to play an equilibrium strategy when other players' preferences are revealed. We discuss a noisy version of the Bayesian Nash equilibrium and a model of strategic ambiguity to account for observed subject behavior.

Keywords: Behavioral Game Theory; Epistemic Game Theory; Nash Equilibrium, Games of Incomplete Information

JEL classifications: C91, C72

[†]Acknowledgements: We would like to thank Peter Dürsch, Jürgen Eichberger, Paul J. Healy, Jörg Oechssler, Stefan Trautmann and Christoph Vanberg for fruitful discussions and helpful comments and references. We would also like to extend our gratitude to the participants of the 8th HeiKaMaX experimental economics workshop, and the ESA European Meeting 2014.

[‡]University of Heidelberg, Department of Economics, Bergheimer Strasse 58, 69115 Heidelberg. Email: cbrunner@uni-hd.de, florian.kauffeldt@awi.uni-heidelberg.de, hannes.rau@gmx.net.

1 Introduction

Applied game theory often relies on the standard Nash equilibrium (Nash 1950, 1951). At the same time, it remains largely unclear which conditions are crucial to ensure that the Nash prediction is accurate. Empirically identifying these conditions is essential to understand in which environments the Nash equilibrium is an appropriate prediction concept. Furthermore, it provides a starting point to develop reasonable alternative concepts for situations where Nash equilibrium fails to accurately predict behavior.

This study is a first attempt to assess the importance of the assumption that agents' preferences are mutually known. We examine this assumption in dominance-solvable normal-form 2×2 games. Previous research suggests that preferences cannot always be assumed to be mutually known. For example, Healy (2011) finds that subjects fail to accurately predict other subjects' preferences over possible outcomes in normal-form 2×2 games. In our experiment, we find that subjects are indeed significantly more likely to play a Nash equilibrium strategy when they are informed about their opponents' preferences over the possible outcomes of the game. Moreover, when preferences are not mutually known, the frequency of equilibrium play is very low in our experiment. Whenever it is unlikely that players know each others' preferences, it might therefore be advisable to use a more general equilibrium concept such as Bayesian Nash equilibrium (Harsanyi, 1967-1968) rather than the standard Nash equilibrium (Nash 1950, 1951).

The experiment consists of two treatments, called "baseline" and "info". Both treatments have two stages. In stage 1, we let subjects rank eight monetary payoff pairs (they will be referred to as "payment pairs"). The first element of such a payment pair corresponds to the amount of money paid to the decision maker. The

second element is the payment that some other subject receives. The same payment pairs are then used to construct four different 2×2 games, or more precisely four different game-forms. In stage 2, each subject plays each of these games exactly once. The two treatments differ in that the preferences elicited in stage 1 are only revealed in stage 2 of treatment info.

This design allows us to avoid the assumption that subjects only care about their *own* monetary payments. Instead, we can use the preferences elicited in stage 1 to describe the game that our subjects play.¹ This is illustrated in Example 1 below, which corresponds to one of the games played in the experiment.

Example 1. *Consider the prisoner’s-dilemma-type game-form in Figure 1. The numbers in the matrix correspond to the amount of money paid to the players, where the first number is the row player’s payment and the second number is the column player’s payment. Suppose that the two players, $i \in \{r, c\}$, where r stands for row and c for column,*

(a) are selfish payment maximizers and only care about their own payments. That is, each player’s preferences over payment pairs $(x_r, x_c) \in \mathbb{R}^2$ are represented by a strictly monotone increasing utility function $v_i(x_i)$ that depends only on his own payment or

(b) have other-regarding preferences represented by a function $\tilde{v}_i : \mathbb{R}^2 \rightarrow \mathbb{R}$,

then the games that result in cases (a) and (b) are depicted in Figure 2.

In Example 1, the game that results if players are selfish (a) is a prisoner’s-dilemma-type game. For all strictly monotone increasing utility functions, v_i , the game has only one Nash equilibrium (U, L) , i.e., everyone defects. That is not necessarily true

¹We maintain the assumption that preferences depend only on players’ monetary payments. That is, the specific game-form, other subjects’ preferences, or any other factors have no effect on subjects’ ordinal ranking of payment pairs. We will discuss evidence suggesting that such considerations do not play an important role in the games used in this study (see Section 3.3).

	L	R
U	4, 4	8, 3
D	3, 8	7, 7

Figure 1: Prisoner’s-dilemma-type game-form

for the induced game (b), where players have social preferences. For example, if $\tilde{v}_r(7, 7) > \tilde{v}_r(8, 3)$ and $\tilde{v}_c(7, 7) > \tilde{v}_c(3, 8)$, then mutual cooperation, (D, R) , is a Nash equilibrium in (b).

	L	R
U	$v_r(4), v_c(4)$	$v_r(8), v_c(3)$
D	$v_r(3), v_c(8)$	$v_r(7), v_c(7)$

(a) Players with selfish preferences

	L	R
U	$\tilde{v}_r(4, 4), \tilde{v}_c(4, 4)$	$\tilde{v}_r(8, 3), \tilde{v}_c(8, 3)$
D	$\tilde{v}_r(3, 8), \tilde{v}_c(3, 8)$	$\tilde{v}_r(7, 7), \tilde{v}_c(7, 7)$

(b) Players with social preferences

Figure 2: Induced games in Example 1

In our experiment, we only ask subjects to rank payment pairs ordinally.² As a result, we cannot compute Nash equilibria in mixed strategies for the induced game (b). In this paper, whenever we refer to a “Nash equilibrium”, we refer to the Nash equilibrium of the induced game using the preferences elicited in stage 1 of the experiment. We focus on those situations where the induced game is dominance-solvable. Theoretically, in such 2×2 games, mutual knowledge of payoff functions along with mutual knowledge of rationality suffices to ensure that agents will play a Nash equilibrium strategy.³ To see this, suppose one player (called “D”) has a strictly

²Eliciting a cardinal ranking of payment pairs would require a more complicated procedure that some subjects might fail to understand. Moreover, it is not obvious that subjects can reliably assign a cardinal utility to each payment pair.

³More generally, Aumann and Brandenburger (1995) show that in normal-form 2×2 games, mutual knowledge of payoff functions, rationality, and conjectures is sufficient to ensure that these conjectures constitute a Nash equilibrium. Conjectures are beliefs about the pure strategy that the other player plays.

dominant strategy. Given that D is assumed to know his own payoff function and is rational, D will play his dominant strategy. The other player (called "ND") knows that D is rational and that he has a strictly dominant strategy. Therefore, ND knows that D will play his dominant strategy. Since ND is himself assumed to be rational and to know his own payoff function, ND must play a best response, and therefore, both agents play a Nash equilibrium strategy.

Furthermore, we focus on situations where the decision maker has a unique pure Nash equilibrium strategy in the induced game. We can only distinguish subjects who play an equilibrium strategy from those who do not in those situations. Moreover, we will also exclude the decisions of subjects who have a strictly or weakly dominant strategy in the induced game. Information about their opponent's preferences is not necessary for those subjects to compute a best response and as a result, information about the other player's preferences should not be expected to have an effect on behavior.

Whenever one player has a unique equilibrium strategy that is not dominant, the other player must have a dominant strategy. In treatment baseline, subjects cannot be certain that this is indeed the case. For example, suppose the row player in the induced game above (b) is selfish. His pure strategy U is then strictly dominant. A column player who prefers $(7, 7)$ to $(8, 3)$ and $(4, 4)$ to $(3, 8)$ then has a unique equilibrium strategy that is not dominant: L . In treatment baseline, such a column player may not be sure whether row is selfish or not and might therefore occasionally play R rather than L . In treatment info, the column player can see that row has a strictly dominant strategy and might therefore play the unique equilibrium strategy L more often. Intuitively, this logic can explain our main result that subjects are much more likely to play a Nash equilibrium strategy in treatment info compared to treatment baseline. We therefore show that subjects not only fail to accurately

predict other players' preferences as previous evidence already suggests, the lack of such information also significantly affects their behavior.

If players do not know each other's preferences, concepts that are more general than Nash equilibrium might provide a more reliable prediction. In our experiment, we find that a strategy is more likely to be played when it cannot lead to the lowest ranked payment pair (maxmin strategy) or when it can result in the realization of the highest ranked one (maxmax strategy). Intuitively, if a subject is uncertain about the strategy choice of his opponent, then, depending on his attitude towards uncertainty, he will try to avoid the lowest ranked payment pair, or, to reach the highest ranked one. We show that the strategic ambiguity model of Eichberger and Kelsey (2014) can rationalize such strategy choices.⁴ This model allows for optimistic responses to strategic ambiguity. Most other strategic ambiguity models such as those of Lo (1996), Eichberger and Kelsey (2000), and Lehrer (2012) assume ambiguity-averse behavior. While these models can explain maxmin strategy choices, they cannot rationalize maxmax behavior.

Another possibility is to take a Bayesian approach by modeling a situation where preferences are not mutually known as a game of incomplete information and to transform it into a Bayesian game (Harsanyi, 1967-1968). Players with different preferences can be thought of as different types and it is then assumed that the prior distribution of types is commonly known. This approach has been used in various fields.⁵ We show that the behavior observed in our baseline treatment is consistent with a noisy version of Bayesian Nash Equilibrium, which we call *Quasi-Bayesian*

⁴“Ambiguity” refers to a situation where probabilities are imperfectly known. Ellsberg (1961) exemplified that individuals frequently display preferences which are not consistent with probabilistic beliefs when facing ambiguity.

⁵In auction theory, for example, the assumption that all bidders are risk neutral and that this is commonly known has been relaxed. Instead, the prior distribution of risk preferences rather than other bidders' actual risk preferences are assumed to be commonly known (see, e.g., Hu and Zou, 2015).

Nash Equilibrium (QBNE).

The papers closest to ours are Healy (2011) and a recent working paper by Wolff (2014). Healy examines whether the sufficient conditions for Nash equilibrium identified by Aumann and Brandenburger (1995) are satisfied when subjects play normal-form 2×2 games in the laboratory. For that purpose, subjects first chose a strategy and then state their beliefs about behavior and preferences of their opponent. Subjects' own preferences and rationality are also measured. Healy finds that there are only very few instances where all conditions are satisfied. Focusing on mutual knowledge of preferences, he finds that both players correctly predict how their opponent ordinally ranked the payment pairs in only 64% of games played. Healy concludes that "The failure of Nash equilibrium stems in a large part from the failure of subjects to agree on the game they are playing." Since mutual knowledge of preferences is one of three conditions that are together sufficient for Nash equilibrium and since the other two are also not fully satisfied in Healy's experiment, it is difficult to assess the impact of the failure of mutual knowledge of preferences on equilibrium play in isolation. By introducing a treatment in which information about the opponent's preferences is directly revealed, we can identify the impact of mutual knowledge on equilibrium play by holding all other factors constant.

Wolff (2014) studies behavior in three-person sequential public good games. In contrast to our experiment, he does not reveal subjects' preferences over the material outcomes. Instead, he elicits subjects' best-response correspondences to the contributions of the other players. In one of his treatments, these are then revealed to all group members. This information has a much smaller effect on the frequency of equilibrium play compared to the treatment effect in our experiment. Revealing best-response correspondences is apparently not sufficient for subjects to be able to predict how much their opponents will contribute: Wolff measures beliefs about others' contributions to

the public good and finds that subjects tend to overestimate these. As a result, they often fail to play an equilibrium strategy even though their contributions tend to be consistent with their beliefs and their own elicited best-response correspondences.

As opposed to the dominance-solvable 2×2 games that we study, several iterations of alternating best responses are required in Wolff's experiment to compute the Nash equilibrium. Some subjects might not be able to do so. Wolff therefore also runs a treatment in which he provides the Nash equilibrium prediction but this additional information does not lead to significantly more equilibrium play. A subject who cannot compute the Nash equilibrium independently might also not understand why he should play a Nash equilibrium strategy. Others might not be confident that other players play the equilibrium strategy. Compared to our experiment, it therefore seems more difficult for subjects to predict how other subjects behave. When it is unclear how other agents behave, information about their preferences is less useful for the purpose of predicting what strategy they will play, which could explain why information about others' best responses has a relatively small impact on the frequency of equilibrium play in Wolff's experiment.

This paper is organized as follows. The next section describes the experimental design. We then present our results and discuss possible explanations for the observed behavior. In particular, we introduce the Quasi-Bayesian Nash Equilibrium model, and show that the behavior observed in our baseline treatment is consistent with it. Section 4 concludes with a summary. In the Appendix, we prove one of the propositions stated in Section 3.4. Furthermore, we describe the strategic ambiguity model of Eichberger and Kelsey (2014) in detail, and show that playing the maxmin or the maxmax strategy is consistent with the model.

2 Experimental design

Our experiment consists of two treatments with two stages each. In the first stage of both treatments, we elicit subjects' preferences over eight different payment pairs. These payment pairs are then used to construct four different 2×2 games. In stage 2 of each treatment, subjects play each one of these games exactly once. In treatment "info", subjects can see their opponent's ordinal ranking of the four payment pairs used in the current game, whereas in treatment "baseline", this information is not disclosed.

We will now describe stage 1 in more detail, which is identical in both treatments. Subjects are asked to create an ordinal ranking over the following set X of eight payment pairs (x_r, x_c) :

$$X = \{(8, 3), (7, 7), (5, 8), (4, 4), (6, 2), (3, 8), (3, 3), (2, 2)\} \subset \mathbb{R}^2 \quad (1)$$

The first number, x_r , corresponds to the amount of money (in Euros) paid to the decision-maker in the role of a row player. The second number, x_c , is paid to some other subject in the role of a column player (the "recipient").⁶ Subjects are informed that they will not interact with the recipient in any other way in either stage of the experiment.

The order in which the payment pairs appear on the screen was randomly determined beforehand and remains constant in all sessions. Subjects rank the payment pairs by assigning a number between one and eight to each pair, where lower numbers indicate a higher preference. The same number can be assigned to multiple payment

⁶Subjects who were assigned the role of a column player ranked the same payment pairs but the first number corresponds to the other player's payoff. Rewriting X for column players such that the first number corresponds to the column player's payment and the second to the row player's, we obtain $X_{column} = \{(8, 3), (7, 7), (8, 5), (4, 4), (2, 6), (3, 8), (3, 3), (2, 2)\} \subset \mathbb{R}^2$.

pairs, thus allowing for indifference. In treatment info, subjects are told that their rankings would be disclosed to other participants at a later stage of the experiment.⁷ In treatment baseline, we made it clear that this information would not be revealed. We will explain at the end of this section how the elicitation of preferences was incentivized. After subjects confirm their ranking, they proceed to stage 2, in which they play the four one-shot 2×2 games in Figure 3 (all numbers are payments in Euro).

Game 1		<i>L</i>	<i>R</i>			<i>L</i>	<i>R</i>
	<i>U</i>	4, 4	8, 3		<i>U</i>	4, 4	8, 3
	<i>D</i>	3, 8	7, 7		<i>D</i>	3, 3	7, 7

Game 2		<i>L</i>	<i>R</i>			<i>L</i>	<i>R</i>
	<i>U</i>	5, 8	7, 7		<i>U</i>	8, 3	2, 2
	<i>D</i>	6, 2	3, 3		<i>D</i>	7, 7	3, 8

Figure 3: Games in the experiment

These games were selected because we conjectured that social preferences might play some role here. Moreover, they could be constructed using only a few payment pairs and exhibit some diversity with respect to the number of pure strategy Nash equilibria under the assumption that subjects are selfish payment maximizers. At least in some of these games, it also appeared likely that only one subject would have a strictly dominant strategy according to the preferences elicited in stage 1 (which turned out to be the case).

In both treatments, subjects can see how they ranked the four payment pairs of the currently played game. This information is displayed by assigning 1-4 stars to each outcome, where more stars indicate a better outcome. In treatment info, subjects

⁷We will discuss the possibility that subjects might strategically misrepresent their preferences in the results section.

are shown both their own *and* their opponent's ranking in matrix-form (see Figure 4). Just like in the payment matrix, the first entry corresponds to the subject's own ranking while the second entry reveals the opponent's ranking. In treatment baseline, subjects are shown the same rankings matrix but this matrix only contains their own rankings.

Game 1

Payoffs:

	left	right
up	4, 4	8, 3
down	3, 8	7, 7

Rankings:
More stars stand for better payoff pairs.

	left	right
up	** **	**** *
down	* ****	*** **

Your decision:

up
 down

OK

Figure 4: Information screen

All subjects play each game exactly once, each time against a different anonymous opponent. Games are played one after another and feedback about the outcome is only provided at the end of the experiment when subjects are paid, but not while subjects still make decisions.

In both treatments, each subject is paid for exactly one of his decisions, which is randomly selected at the end of the experiment. If a decision from stage 1 is chosen,

two of the eight payment pairs from (1) are randomly selected. The row subject is then paid the first number, x_r , of the payment pair that he ranked more highly in stage 1. The second number, x_c , is paid to some other column subject. Column subjects are paid in a similar manner. The probability that stage 1 is paid is $\frac{7}{8}$ while stage 2 is paid with a probability of $\frac{1}{8}$. These probabilities are consistent with selecting each of the $\binom{8}{2}$ possible pairs of payment pairs and each of the four decisions made in stage 2 with equal probability. Paying stage 1 with a substantially higher probability also reduces the odds that subjects might misrepresent their preferences. This issue will be discussed in more detail in Section 3.3.

Subjects were given printed instructions and could only participate after successfully answering several test questions. Test questions as well as the rest of the experiment were programmed using Z-Tree (Fischbacher, 2007). All sessions were conducted between August and October 2014 at the AWI-Lab of the University of Heidelberg. Subjects from all fields of study were recruited using Orsee (Greiner, 2004). Fewer than half of the subjects were economics students. Sessions lasted about 40-50 minutes on average. The following table summarizes the number of participants per session as well as average payments:

Table 1: Summary of treatment information

Treatment	Sessions	Subjects	Average payment
baseline	8	84	€ 12.36
info	7	80	€ 12.59

3 Results

In this section, we first characterize subjects' preferences as measured in stage 1 of the experiment. We then present the main treatment effect: subjects are significantly more likely to play their unique equilibrium strategy in treatment info than in treatment baseline. This effect can be observed in each of the four games, but it is not significant for every game when we only use the data from one single game at a time. We then discuss the possibility that subjects misrepresent their true preferences and that preferences change when subjects are shown their opponents' preferences and find no evidence for either of these effects. The section concludes with a discussion of two possible ways to analyze a situation in which preferences are not mutually known: A noisy version of Bayesian Nash equilibrium and models of strategic ambiguity.

3.1 Characterization of measured preferences

In stage 1 of the experiment, we elicit subjects' preferences over the payment pairs $(x_r, x_c) \in X \subset \mathbb{R}^2$ defined in equation (1). Tables 2 and 3 show the ordinal rankings stated by at least two subjects who were assigned the role of a row and column player respectively. Payment pairs that are assigned a lower number are preferred to payment pairs with a higher number. A full list of these rankings can be found in the appendix (tables 9 and 10).

Table 2: Preferences stated by at least two subjects who were assigned the role of a row player, both treatments. Smaller numbers are assigned to better ranked payment pairs.

(8,3)	(7,7)	(5,8)	(4,4)	(6,2)	(3,8)	(3,3)	(2,2)	n
1	2	4	5	3	6	7	8	28
1	2	4	5	3	7	6	8	10
1	2	4	5	3	6	6	7	6
2	1	4	5	3	6	7	8	6
1	2	3	5	4	6	7	8	4
2	1	3	5	4	6	7	8	4
2	1	3	4	5	6	7	8	2
1	2	3	4	3	5	6	7	2
2	1	3	6	4	5	7	8	2
1	2	3	4	5	6	7	8	2
1	2	4	6	3	5	7	8	2

Table 3: Preferences stated by at least two subjects who were assigned the role of a column player, both treatments. Smaller numbers are assigned to better ranked payment pairs.

(8,3)	(7,7)	(8,5)	(4,4)	(2,6)	(3,8)	(3,3)	(2,2)	n
2	3	1	4	7	5	6	8	35
3	2	1	4	7	5	6	8	6
3	2	1	5	7	4	6	8	4
3	2	1	5	6	4	7	8	3
3	1	2	4	7	5	6	8	3
3	1	2	4	8	6	5	7	2
1	2	1	3	6	4	5	7	2
1	3	2	4	8	6	5	7	2

To characterize subjects' preferences, we introduce four properties: pareto-efficiency, strict pareto efficiency, maximization of own payoff, and maximization of total payoff. These properties are defined as follows:

Definition 1 (Pareto efficiency). A subject's preferences \succsim on X are said to satisfy *pareto-efficiency* if, for all $x, y \in X$, $x \succ y$ whenever $x_1 \geq y_1$ and $x_2 \geq y_2$ with at least one inequality strict.

Definition 2 (Strict pareto efficiency). A subject's preferences \succsim on X are said to satisfy *strict pareto-efficiency* if, for all $x, y \in X$, $x \succ y$ whenever $x_1 > y_1$ and $x_2 > y_2$.

Definition 3 (Maximization of own payoff). A subject is said to *maximize his own payoff* if, for all $x, y \in X$, $x \succ y$ whenever $x_1 > y_1$.

Definition 4 (Maximization of total payoff). A subject is said to *maximize total payoff* if, for all $x, y \in X$, $x \succ y$ whenever $x_1 + x_2 > y_1 + y_2$.

Table 4 shows the fraction of subjects whose preferences are consistent with the properties defined above.

Table 4: Measured preferences

	Pareto efficiency	Strict pareto efficiency	Maximization of own payoff	Maximization of total payoff	n
Percentage consistent (pooled)	69.5%	92.1%	54.9%	1.8%	164
Percentage in baseline	67.9%	91.7%	50.0%	2.4%	84
Percentage in info	71.3%	92.5%	60.0%	1.3%	80

3.2 Main result

Our main hypothesis is that subject behavior is more consistent with the Nash equilibrium when preferences are mutually known. We test this hypothesis by using two different subsets of our data. Recall that each subject played four games. Since there are a total of 164 subjects who participated in the experiment, we have data on 656 individual decisions, 336 in treatment baseline and 320 in treatment info. We exclude those decisions where both strategies are played with strictly positive probability in some Nash equilibrium, which leaves us with 425 decisions (213 in treatment baseline and 212 in treatment info). We also exclude those decisions where one pure strategy is weakly or strictly dominant. In such a situation, the best response does not depend on the other player’s strategy and therefore, it should not matter whether or not the other players’ preferences are known. This leaves us with 147 individual decisions, 75 in treatment baseline and 72 in treatment info. In all of these 147 games, the subject whose decision we study has a unique pure equilibrium strategy and that subject’s opponent has a strictly dominant strategy. We test our main hypothesis using these 147 observations and will refer to the according subset of our data as “all subjects”. Figure 5 shows that subjects play an equilibrium strategy more often in treatment info than in treatment baseline.

To test whether these differences are significant, we run a logit regression. The dependent variable “equilibrium strategy played” assumes a value of 1 if a subject plays the unique equilibrium strategy and 0 otherwise. We include an intercept as well as a dummy variable, which assumes a value of 1 if the observation is generated in treatment info and 0 otherwise. These results are shown in Table 5. The treatment effect is highly significant. Informing subjects about their opponents’ preferences leads to a significantly higher frequency of equilibrium play.




Figure 5: Frequencies of played unique equilibrium strategies

Table 5: Logit regression “equilibrium strategy played”, robust standard errors clustered by subject

Dependent variable: equilibrium strategy played	All Subjects		Consistent subjects only	
	Coefficient	SE	Coefficient	SE
info	0.93***	0.36	1.37***	0.43
constant	-0.88***	0.26	-1.15***	0.32
n		147		108
Clusters		109		76
Pseudo R^2		0.038		0.079

*** significant at 1% level

We run the same test a second time with a smaller subset of our data which no longer includes the decisions made by subjects who played a strictly dominated strategy in at least one of the four games. Either the preferences that these subjects stated in stage 1 do not reflect their true preferences or they are not rational in the sense that their choice in stage 2 is inconsistent with their stated preferences. Table 6 shows that approximately one third of our subjects violate strict dominance at least once.

Table 6: Violations of strict dominance

Treatment	Subjects	Games played	Games with dominated strategy	Dominated strategy played	Subjects who played dominated strategy at least once
Baseline	84	336	136	23.53%	32.14%
Info	80	320	140	25.71%	33.75%

Just like in subset “all subjects” we also only use games where the subject has a unique equilibrium strategy that is not dominant. Removing the choices made by inconsistent subjects therefore further reduces the number of observations to 108 individual decisions, 54 in treatment baseline and 54 in treatment info. We will refer to this subset of our data as “consistent subjects only”. The treatment effect is even stronger when we use these consistent subjects only. We also test whether there is a significant treatment effect using a two-tailed two-sample Wilcoxon rank-sum test. The dependent variable is the frequency with which a subject played an equilibrium strategy. Each subject who plays at least one game where the subject has a unique equilibrium strategy that is not weakly dominant counts as one observation. We run

the same test for all subjects and for consistent subjects only. When using all (only consistent) subjects, we have 54 (38) observations in treatment baseline and 55 (38) in treatment info. We can reject the null hypothesis that the distribution of the frequency of equilibrium play is the same in both treatments regardless of which data set we use ($p=0.038$ using all subjects, $p=0.008$ using consistent subjects only).

Result 1. *Subjects are more likely to play their unique Nash equilibrium strategy when preferences are mutually known.*

As a robustness check, we also compute the frequency of equilibrium play for each game separately. These results are shown in Figure 6 for all subjects and in Figure 7 for consistent subjects only. Regardless of which subset of our data we use, the frequency of equilibrium play is higher in treatment info than in treatment baseline for every single game. However, using a Fisher exact test, this difference is only significant at the 5% level for game 3, regardless of whether we use all subjects or consistent subjects only. We have more observations for game 3 than for any other game. In game 3, it occurred particularly often that one subject had a strictly dominant strategy while the other subject did not have a strictly or weakly dominant strategy. Details of these tests can be found in the appendix (tables 12 and 13).

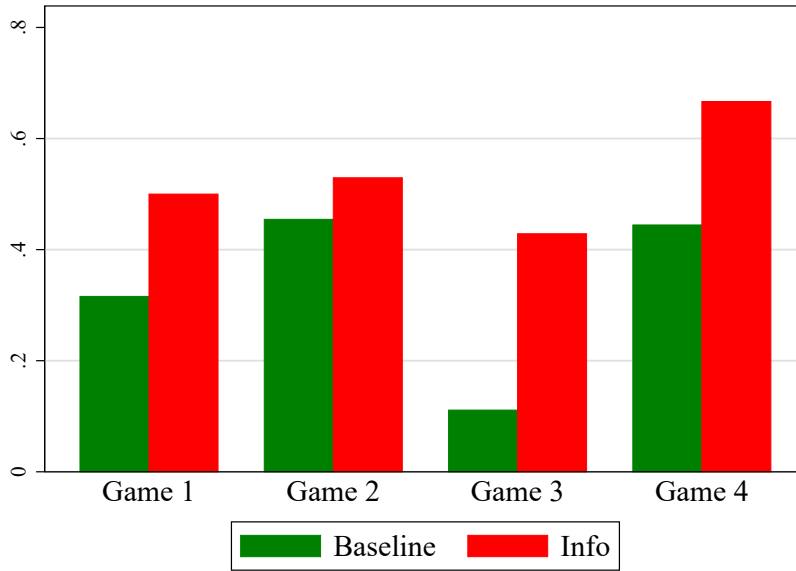


Figure 6: Frequency of equilibrium play by game, all subjects

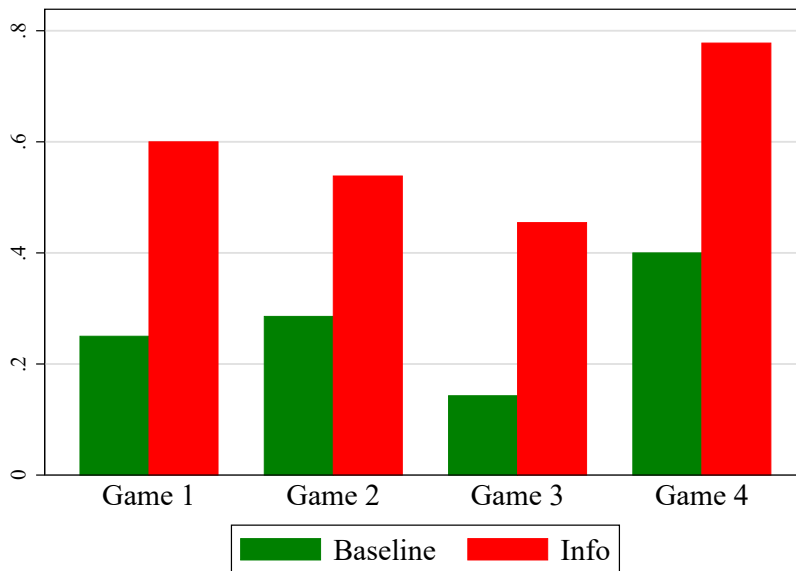


Figure 7: Frequency of equilibrium play by game, consistent subjects only

3.3 Did we manage to elicit subjects' true preferences?

When preferences are elicited in stage 1 of the experiment, subjects in treatment info are aware that these preferences will be revealed to other subjects. However, they are not informed about the specific games that are played in stage 2. Hence, subjects did not have the information necessary to figure out what kind of misrepresentation might be most advantageous: in some 2×2 games, it could be beneficial to be perceived as having social preferences whereas in other games, the contrary is more likely (e.g., in the chicken game). Moreover, recall that a decision made in stage 2 affects a subject's payment with a probability of only $1/8$. Therefore, it does not seem plausible that a rational subject would misrepresent his preferences in stage 1.

We test the claim that subjects truthfully state their preferences in stage 1 of treatment info by using the frequency with which subjects play strictly dominated strategies in stage 2 of the experiment. To identify strategies that are strictly dominated, we use the preferences elicited in stage 1. If these reflect a subject's true preferences, a rational subject should never play such a strictly dominated strategy. In contrast, if subjects strategically misrepresent their preferences in stage 1, a strategy that we classify as strictly dominated may in fact not be dominated according to the subjects' true preferences. Since preferences in treatment baseline are not revealed to other subjects, it is clear that subjects in treatment baseline have no reason to misrepresent their preferences. Therefore, we can compare the frequency with which subjects play a strictly dominated strategy in the two treatments to test the claim that preferences are truthfully revealed in stage 1 of treatment info. If that claim is true, no difference should be observed. Otherwise, subjects should be more likely to play a strictly dominated strategy in treatment info than in baseline.

Table 6 shows how often subjects play a strictly dominated strategy using the

preferences stated in stage 1 to define the according games. Each subject played 4 games, thus resulting in 336 games played in treatment baseline and 320 in info. In 136 of these games in treatment baseline and 140 in info, one of the strategies was strictly dominated. In roughly a quarter of these cases, the strictly dominated strategy was played.

In order to check the assumption that subjects do not misrepresent their preferences in both treatments, we run a logit regression using the 136 games in treatment baseline as well as the 140 games in treatment info as observations. The dependent variable “dominated strategy played” is a dummy variable that assumes a value of 1 if the strictly dominated strategy was played. The only explanatory variable other than the intercept is a treatment dummy (“info”) (see Table 7).

Table 7: Logit regression “dominated strategy played”, robust standard errors clustered by subject

Dependent variable: dominated strategy played		
	Coefficient	SE
Info	0.12	0.30
Constant	-1.18***	0.22
n		276
Clusters		160
Pseudo R^2		0.0006

*** significant at 1% level

The coefficient estimate for the treatment dummy is not significantly different from 0. Hence, the null hypothesis cannot be rejected. We also test the same assumption using a two-tailed two-sample Wilcoxon rank-sum test. The dependent variable is then the frequency with which a subject plays a dominated strategy. Each subject who had a strictly dominant strategy in at least one of the four games corresponds to an observation. There are 81 such observations in treatment baseline and 79

in treatment info. We cannot reject the null hypothesis that the frequency with which strictly dominated strategies are played follows the same distribution in the two treatments ($p=0.93$).

Result 2. *Subjects are equally likely to play a strictly dominated strategy in both treatments.*

Moreover, the fraction of subjects whose reported preferences are consistent with own payoff maximization is even slightly larger in treatment info compared to treatment baseline, though the difference is not significant ($p= 0.21$ using a Fisher exact test). We therefore maintain the assumption that subjects truthfully state their preferences in stage 1 of the experiment in both treatments.

In psychological game theory, Rabin (1993) and Dufwenberg and Kirchsteiger (2004) introduced models of reciprocity in which players reward kind actions and punish unkind ones. Reciprocity could lead to a problem equivalent to the misrepresentation of preferences discussed in this section. For instance, consider game 1 in stage 2 of treatment info. Suppose an own-payoff maximizer (row) is matched with a total-payoff maximizer (column). The row player might then believe that column will cooperate (play R), even though column expects row to defect (play U). This expected kindness on the part of column might then induce row to also cooperate, thus violating our assumption that only outcomes matter. In other words, subjects' preferences might change once they are shown their opponents' ranking of payment pairs in stage 2 of treatment info. If so, the preferences over outcomes we use in our analysis would no longer correspond to subjects' true preferences, or at least not govern subjects' strategic behavior. Since such preference adjustments are only possible in treatment info but not in treatment baseline, we can use Result 2 to argue that such effects probably do not matter much in our experiment. If they did, one

would expect to observe that subjects play strictly dominated strategies more often in treatment info compared to treatment baseline.

3.4 Possible explanations for observed subject behavior

Our main results show that making sure that payoffs are mutually known has a significant effect on behavior. Moreover, subjects rarely play their unique non-dominant equilibrium strategy when other player's preferences are not revealed. Whenever payoffs are unlikely to be mutually known, it might therefore be advisable to use a more general model than the standard Nash equilibrium. In this section, we will first discuss a noisy version of the Bayesian Nash Equilibrium, which we call Quasi-Bayesian Nash Equilibrium (QBNE). Subsequently, we describe the strategic ambiguity model of Eichberger and Kelsey (2014), and show that this concept can rationalize the maxmin and maxmax strategy choices that we frequently observe in our experiment.

3.4.1 A quasi Bayesian approach

As described in the introduction, the games played in the baseline treatment can also be considered as Bayesian games in which players with different preferences represent different types. It is then assumed that the prior distribution of types is commonly known. In what follows, we first introduce our concept of a Quasi-Bayesian Nash equilibrium (QBNE). Subsequently, we describe in detail how to model the situation in the baseline treatment of our experiment as a Bayesian game. Finally, we show that the predictions of the Quasi-Bayesian Nash Equilibrium model are not trivially consistent with our data and could thus be falsified. It is therefore interesting to observe that the data collected in the baseline treatment is indeed consistent with such a model. However, it should be noted that the predictions of QBNE are not

precise, particularly given that we elicit an ordinal rather than a cardinal ranking of payment pairs in stage 1 of the experiment.

In a Bayesian game, a strategy σ_i of player i prescribes a mixed action for each possible type of player i , formally $\sigma_i : \Theta_i \rightarrow \Delta(A_i)$, where Θ_i denotes player i 's type space and $\Delta(A_i)$ the set of mixed actions of player i (i.e., the set of all probability distributions over player i 's set of pure actions A_i). Denote by Σ_i the set of all strategies of player i and let $\Sigma = \prod_{i \in I} \Sigma_i$. The interim expected utility of player i with type $\theta_i \in \Theta_i$ given a mixed strategy profile $\sigma \in \Sigma$ is formed by taking the expectation with respect to the mixed strategy profile σ and the expectation with respect to the conditional type distribution, formally,

$$EU_i(\sigma | \theta_i) = \sum_{\theta_{-i} \in \Theta_{-i}} \pi(\theta_{-i} | \theta_i) \left(\sum_{a \in A} \left(\prod_{j \in I} \sigma_j(a_j | \theta_j) \right) u_i(a, \theta_i, \theta_{-i}) \right),$$

where $\pi(\theta_{-i} | \theta_i)$ denotes the probability of θ_{-i} under the condition that i knows he is of type θ_i , and $\sigma_j(a_j | \theta_j)$ is the probability of action a_j that strategy σ_j prescribes for θ_j . As usual “ $-i$ ” stands for “all players except player i ”.

Recall that a substantial fraction of subjects played strategies that are strictly dominated and thus inconsistent with the preferences stated in the first stage (see Table 6). To account for such irrational behavior, we add a noisy type $\tilde{\theta}_i$ to each player's type space that randomly selects a pure strategy. In a QBNE all types other than this noisy type play a best response given the commonly known distribution of types, while the noisy type plays an arbitrary proper mixed action.

Definition 5. A *Quasi-Bayesian Nash equilibrium (QBNE)* for a Bayesian game is a strategy profile $(\sigma_i^*, \sigma_{-i}^*) \in \Sigma$ such that, for each player $i \in I$,

- (i) $\sigma_i^*(\theta_i) \in \arg \max_{\sigma_i \in \Sigma_i} EU_i(\sigma_i, \sigma_{-i}^* | \theta_i)$ for all non-noisy types $\theta_i \in \Theta_i \setminus \{\tilde{\theta}_i\}$, and

- (ii) $\sigma_i^*(\tilde{\theta}_i) \in \text{int}(\Delta(A_i))$ for the noisy type $\tilde{\theta}_i \in \Theta_i$, where $\text{int}(\Delta(A_i))$ denotes the interior of the set of player i 's mixed actions.

Obviously, a QBNE is weaker than a Bayesian Nash equilibrium since it only requires that non-noisy types play mutual best responses. The existence of a QBNE follows from the standard fixed-point argument by Nash (1950, 1951).

Modeling the situation in the baseline treatment of our experiment as a Bayesian game

In order to show that each game in stage 2 of the baseline treatment can be considered as a Bayesian game, we make the following conceptual distinction:

- *A subject* is an individual who took part in the baseline treatment of our experiment. Hence, in total, there are 84 subjects. We denote one subject by $k \in \{1, \dots, 84\}$.
- *Player* is the role or the position of a subject in one of the four 2×2 games in the baseline treatment. There are two types of players: row and column. Henceforth, given one of the four games, the subject who is the row player is denoted by r and the column player by c . We denote one player by $i \in \{r, c\}$.
- *Row and column types* characterize potential characteristics of row and column players. *Non-noisy types* play a best response given their preferences and the commonly known distribution of types. Each non-noisy type is characterized by an ordinal preference ordering over the eight payment pairs from equation (1). In addition, there is a *noisy type* that randomly selects a pure strategy.

We will use the observed fraction of row (column) subjects who played a strictly dominated action at least once (see Table 6) as an estimator for the probability that

a row (column) player is of the noisy type.⁸

Assumption 1. *The fraction of row and column subjects who played a strictly dominated action at least once corresponds to the probability of noisy types of row and column players.*

At the first stage of treatment baseline, we elicit each subject k 's ordinal preferences \succsim_k over eight payment pairs.

Definition 6. Subject k 's ordinal preference ordering \succsim_k on the set X defined in equation (1) is a function $f_k : X \rightarrow \{1, \dots, 8\}$.

We do not know subjects k 's utility function $v_k(\cdot)$ exactly, but we know that $v_k(\cdot)$ is a representation of the ordinal ordering \succsim_k , i.e., for all $x, y \in X$ and all k , $v_k(x) \geq v_k(y)$ if and only if $x \succsim_k y$. Since each non-noisy type corresponds to a specific ordinal ranking, we shall assume that the utility functions of all non-noisy subjects who stated the same ordinal ranking are identical:

Assumption 2. *For any two subjects k and k' , who are either both row or both column players and who never played a strictly dominated action, if $\succsim_k = \succsim_{k'}$, then $v_k(\cdot) = v_{k'}(\cdot)$.*

Again, we take the observed relative frequencies as estimators for the probabilities of a non-noisy type. For instance, the probability of the non-noisy row type $\theta_r^{\succsim'}$ that is characterized by the ordinal ranking \succsim' equals

$$\pi[\theta_r^{\succsim}] = \frac{\# \text{ non-noisy row subjects with ordinal ranking } \succsim'}{\# \text{ row subjects}}. \quad (2)$$

Following Harsanyi (1967-1968), we assume that the types and the prior distribution of types, i.e., the probabilities, are commonly known.

⁸Note that this is a biased estimator since noisy subjects may accidentally behave consistently. That is, the estimator systematically underestimates the probability of the noisy type.

Assumption 3. *Row and column players' beliefs over types, π_r and π_c , correspond to the relative frequencies of types in the experiment.*

Assumption 3 implies that a strictly positive probability is assigned to all types that we observed in the baseline treatment, whereas all other types occur with zero probability.

Consistency between observed behavior and Quasi-Bayesian Nash equilibrium

This section provides two propositions. The first one shows that the predictions of Quasi-Bayesian Nash equilibrium are falsifiable using our data. In the second proposition, we show that the behavior observed in the baseline treatment is consistent with a QBNE. Both propositions are based on the following notion of consistency between action choices and QBNE: we say that a QBNE is consistent with an action set combination if the action set combination lies in the support of the QBNE. The *support of a QBNE* is the set of all action profiles that are played with strictly positive probability.

In the following, we will formally define this consistency notion. In order to do so, we first introduce some notation and definitions that will be used throughout this section. Suppose that in each of the four interactive situations in stage 2, the subjects played a Bayesian game according to Assumptions 1, 2, and 3. Let Θ_i be the set of types of player $i \in \{r, c\}$ that have non-zero probability. From now on, when speaking of “types”, we refer to types that have non-zero probability. An *action set combination* a^G in game $G \in \{1, 2, 3, 4\}$ is an ordered set that assigns at least one action from $A_r = \{U, D\}$ to each row type and at least one action from $A_c = \{L, R\}$ to each column type.

Definition 7. An action set combination a^G in game $G \in \{1, 2, 3, 4\}$ is a set

$$a^G \subseteq \{U, D\}^{|\Theta_r|} \times \{L, R\}^{|\Theta_c|},$$

where $|\Theta_r|$ and $|\Theta_c|$ denote the number of different row types and column types.

Let $\sigma_i(\theta_i)$ be the mixed action that strategy σ_i prescribes for type θ_i . Denote by $\sigma_i(a_i | \theta_i)$ the probability with which θ_i plays action a_i according to $\sigma_i(\theta_i)$. Given a row or column type θ_i , an action $a_i \in A_i$ is said to be contained in the *support of* $\sigma_i(\theta_i)$ if type θ_i plays a_i with strictly positive probability, formally $\text{supp}(\sigma_i(\theta_i)) = \{a_i \in A_i | \sigma_i(a_i | \theta_i) > 0\}$. The *support of a type-contingent strategy* of row or column, σ_i , equals the Cartesian product of the supports of the strategy for all given types:

$$\text{supp}(\sigma_i) = \prod_{\theta_i \in \Theta_i} \text{supp}(\sigma_i(\theta_i)).$$

Finally, the *support of a type-contingent strategy profile* $(\sigma_r, \sigma_c) \in \Sigma_r \times \Sigma_c$ is the Cartesian product of the supports of the type-contingent strategies σ_r and σ_c .

Definition 8. The *support of a strategy profile* $(\sigma_r, \sigma_c) \in \Sigma_r \times \Sigma_c$ is the set

$$\text{supp}(\sigma_r, \sigma_c) = \text{supp}(\sigma_r) \times \text{supp}(\sigma_c) \subseteq \{U, D\}^{|\Theta_r|} \times \{L, R\}^{|\Theta_c|}.$$

Now, we are ready to formally define our notion of consistency between an action set combination and a Quasi-Bayesian Nash equilibrium.

Definition 9. An action set combination a^G in game G is said to be consistent with a QBNE $\sigma^* \in \Sigma_r \times \Sigma_c$ of the Bayesian game that results from G if $a^G \subseteq \text{supp}(\sigma^*)$.

Our first proposition shows that the predictions of QBNE can be falsified using our data and Definition 9. The reason is that some of the eight payment pairs appear in more than one of the four games in the second stage of the experiment. This can lead to a contradiction as the proof of the proposition illustrates.

Proposition 1. *Under Assumption 1, 2, and 3, there exist action set combinations $a^1, a^3 \subseteq \{U, D\}^{|\Theta_r|} \times \{L, R\}^{|\Theta_c|}$ such that, if a^1 is consistent with a QBNE of the Bayesian game that results from game 1, then a^3 is not consistent with any QBNE of the Bayesian game that results from game 3, and vice versa.*

Proof. Consider games 1 and 3 and the column types θ_7 and θ_8 (see Table 10). Suppose that both types have the action set $\{L, R\}$ in both games. Action set combinations that involve this action set for θ_7 and θ_8 are consistent with Quasi-Bayesian Nash equilibria for the Bayesian games that result from game 1 and 3 only if the types θ_7 and θ_8 are indifferent between the actions L and R in both games, given the equilibrium strategy of row. Let v_7, v_8 be utility functions that represent the ordinal rankings of θ_7 and θ_8 . Then, type θ_7 is indifferent between his actions in game 1 if he expects row to play U with probability

$$\frac{v_7(7, 7) - v_7(3, 8)}{(v_7(4, 4) - v_7(8, 3) + v_7(7, 7) - v_7(3, 8))} \quad (3)$$

and in game 3 if he expects row to play U with probability

$$\frac{v_7(7, 7) - v_7(3, 3)}{(v_7(4, 4) - v_7(8, 3) + v_7(7, 7) - v_7(3, 3))}. \quad (4)$$

Similarly, type θ_8 is indifferent between his actions in games 1 and 3 if he expects row to play U with probabilities

$$\frac{v_8(7, 7) - v_8(3, 8)}{(v_8(4, 4) - v_8(8, 3) + v_8(7, 7) - v_8(3, 8))} \quad \text{and} \quad (5)$$

$$\frac{v_8(7, 7) - v_8(3, 3)}{(v_8(4, 4) - v_8(8, 3) + v_8(7, 7) - v_8(3, 3))}. \quad (6)$$

Furthermore, θ_7 and θ_8 have to have the same expectation about the probability

with which row plays U . The reason is that, under Assumptions 1, 2, and 3, given a strategy σ'_r of row, each column type expects his opponent to play U with probability:

$$\beta[U | \sigma'_r] = \sum_{\theta_r \in \Theta_r} \pi[\theta_r] \cdot \sigma'_r(U | \theta_r).$$

where $\pi[\theta_r]$ denotes the probability of the row type θ_r (i.e., according to Assumption 3, the relative frequency of θ_r types), and $\sigma'_r(U | \theta_r)$ the probability with which type θ_r plays action U according to σ'_r .

Consequently, action set combinations in game 1 and 3, where the action set of θ_7 and θ_8 equals $\{L, R\}$, are simultaneously consistent with some QBNE in the Bayesian games 1 and 3 only if

- (i) θ_7 's expectation about the probability with which row plays U equals expression (3) in game 1 and (4) in game 3, and, θ_8 's expectation equals (5) in game 1 and expression (6) in game 3.
- (ii) θ_7 and θ_8 have the same expectations in both games, i.e., it holds that (3)=(5), and, (4)=(6).

According to the ordinal ranking that characterizes θ_7 , it holds that $v_7(8, 3) > v_7(3, 3)$. This implies that (3) < (4), i.e.,

$$\frac{v_7(7,7) - v_7(3,8)}{(v_7(4,4) - v_7(8,3) + v_7(7,7) - v_7(3,8))} < \frac{v_7(7,7) - v_7(3,3)}{(v_7(4,4) - v_7(8,3) + v_7(7,7) - v_7(3,3))}.$$

Similarly, for θ_8 , it holds that $v_8(8, 3) < v_8(3, 3)$, which implies (5) > (6), i.e.,

$$\frac{v_8(7,7) - v_8(3,8)}{(v_8(4,4) - v_8(8,3) + v_8(7,7) - v_8(3,8))} > \frac{v_8(7,7) - v_8(3,3)}{(v_8(4,4) - v_8(8,3) + v_8(7,7) - v_8(3,3))}.$$

These inequalities taken together contradict the condition that (3)=(5) and, at the same time, (4)=(6). Hence, if the action set combination in game 1 is consistent with some QBNE in Bayesian game 1, then the action set combination in game 3 is not consistent with any QBNE for the Bayesian game 3, and vice versa. \square

Our second proposition shows that the action set combinations \hat{a}^G that we observed in the experiment in the games $G \in \{1, 2, 3, 4\}$ are consistent with the Quasi-Bayesian Nash equilibrium model in the sense of Definition 9. Before we state the proposition, we need to specify what we mean by observed action sets. Let a_k^G be the action played by subject k in game $G \in \{1, 2, 3, 4\}$ in the second stage of the baseline treatment. *The observed actions played in game G that are associated with a given row or column type θ_i* are denoted by $a_{\theta_i}^G$. These action sets are defined as the union of all action choices of subjects who are of type θ_i , formally

$$a_{\theta_i}^G = \bigcup_{k \text{ is of type } \theta_i \in \Theta_i} a_k^G \subseteq A_i.$$

Furthermore, *the observed action sets in game G that are associated with all types of row and column respectively Θ_i , $i \in \{r, c\}$* , are

$$a_{\Theta_i}^G = \bigtimes_{\theta_i \in \Theta_i} a_{\theta_i}^G \subseteq A_i^{|\Theta_i|}.$$

Finally, *the observed action set combination \hat{a}^G in game G* is

$$\hat{a}^G = a_{\Theta_r}^G \times a_{\Theta_c}^G.$$

Proposition 2. *Under Assumption 1, 2, and 3, there exists a QBNE σ^* for every Bayesian game G such that the observed action set combination \hat{a}^G is consistent with σ^* for all $G \in \{1, 2, 3, 4\}$.*

Proof. See appendix.

3.4.2 A non-Bayesian approach

While a Bayesian approach requires common knowledge of the distribution of types, subjects in our experiment might not hold probabilistic beliefs about other

players' types or choice of strategy. Depending on subjects' attitude towards uncertainty, it can then seem reasonable for a pessimistic agent to play a strategy that cannot lead to the lowest ranked payment pair (maxmin). Lo (1996), Eichberger and Kelsey (2000), and Lehrer (2012) are examples of models of strategic ambiguity that are consistent with such behavior. Alternatively, agents might try to reach the highest ranked payment pair (play a maxmax strategy). In the appendix, we show that the strategic ambiguity model of Eichberger and Kelsey (2014) is consistent with playing either a maxmin or a maxmax strategy though that model also does not provide a precise prediction for the games discussed here.

Playing a maxmin or a maxmax strategy can be a response to uncertainty about other players' payoff functions. It can also be a response to uncertainty about whether the opponent is rational. In treatment baseline, our subjects face both types of uncertainty while the uncertainty about other players' payoffs is removed in treatment info. Since there is some uncertainty in both treatments, we would expect a strategy to be played more often if it cannot lead to the lowest ranked payment pair (maxmin strategy) or if it can lead to the highest ranked payment pair (maxmax strategy) in both treatments. Both effects are expected to be stronger in treatment baseline compared to treatment info.

We test these conjectures by running a conditional logit regression. An observation corresponds to a pure strategy. The dependent variable ("played") assumes a value of 1 if a strategy is played and 0 otherwise. Three independent variables are used to characterize each strategy: "equilibrium" indicates whether a strategy is a Nash equilibrium strategy. "maxmax" assumes a value of 1 if a strategy can lead to a most highly ranked payment pair. "maxmin" indicates whether a strategy can result in the realization of a lowest ranked payment pair (maxmin = 0 if that is the case, maxmin = 1 otherwise). We only use decisions made by subjects who never played a

strictly dominated strategy. Table 8 shows that whether or not a strategy is a Nash equilibrium strategy only matters in treatment info when predicting which strategies subjects will play. In contrast, the coefficients of maxmax and maxmin are significant in both treatments.

Table 8: Conditional logit regression “played”, robust standard errors clustered by subject

Dependent variable: played	Baseline		Info	
	Coefficient	SE	Coefficient	SE
equilibrium	-0.25	0.35	1.02***	0.34
maxmax	1.60***	0.31	1.07***	0.22
maxmin	1.53***	0.29	1.30***	0.21
n		456		424
Clusters		57		53
Pseudo R^2		0.42		0.42

*** significant at 1% level

Result 3. *In both treatments, a strategy is more likely to be played when it cannot lead to the lowest ranked payment pair and when it can lead to the highest ranked payment pair.*

The coefficient estimate for the variable “equilibrium” differs significantly among the two treatments and is only useful to predict play in treatment info but not in treatment baseline. In contrast, the highest and lowest ranked payment pair seems to attract our subjects’ attention in both treatments. As expected, the according coefficient estimates are higher in treatment baseline than in treatment info. However, the difference is not significant.⁹

⁹The coefficient estimate of an interaction term of maxmin and the treatment dummy (maxmax and the treatment dummy) is not significant at the 5% level.

4 Conclusion

The assumption that payoffs are mutually known is often not satisfied in the laboratory. It seems plausible that similar difficulties exist in many real-world situations as well. Our experiment shows that it is a relevant assumption: making sure that payoffs are mutually known leads to significantly more equilibrium play.

When deciding what model to apply to a specific situation, whether or not agents can reasonably be expected to know other agents' payoff functions should therefore play an important role. At least in the simple 2×2 games we analyzed, subjects are very unlikely to play a Nash equilibrium strategy when payoffs are not mutually known. It might then be worthwhile to apply a more complex model such as a strategic ambiguity model or the Bayesian Nash equilibrium of Harsanyi (1967-1968), even though such models tend to provide less precise predictions. Both the strategic ambiguity concept of Eichberger and Kelsey (2014) and a noisy version of the Bayesian Nash equilibrium are consistent with the data of our baseline treatment where preferences are unlikely to be mutually known.

A Appendix

A.1 Proof of Proposition 2

Proof. Suppose that row subjects' expectation that their opponent will play L in G equals $\beta_r^G \in (0, 1)$ and column players' expectation that their opponent will play U in G is $\beta_c^G \in (0, 1)$. Let \succsim_i be the ordinal ranking that is associated with type $\theta_i \in \Theta_i$.

The proof is organized as follows. In Lemma 1, we show that there exists a utility function v_i for all non-noisy types $\theta_i \in \Theta_i$, which represents \succsim_i , such that if $\hat{a}_{\theta_i}^G = \{U, D\}$ (resp. $\hat{a}_{\theta_i}^G = \{L, R\}$), then θ_i is indifferent between his pure actions, given the beliefs β_r^G and β_c^G respectively. Subsequently, by using Lemma 1, we prove that there exists a QBNE σ_G^* for the Bayesian game of $G \in \{1, 2, 3, 4\}$ such that if $\hat{a}_{\theta_i}^G = \{U, D\}$ (resp. $\hat{a}_{\theta_i}^G = \{L, R\}$), then $\sigma_G^*(\theta_i)$ prescribes a proper mixed action for each non-noisy type θ_i .

Note that we do not have to consider noisy types, and non-noisy types who have a strictly dominant action. The played actions of both types are always consistent with a QBNE. Furthermore, the proof is trivial for non-noisy types, who have in only one of the four games no strictly dominant action. In tables 9 and 10 in the next section, the remaining types are labeled through a type name. Furthermore, Table 11 shows which games are relevant and the action sets associated with each of the remaining types in every game.

Lemma 1. *Consider the types θ_j , $j = 1, \dots, 11$, defined labeled in tables 9 and 10. Given game G , let $\beta_r^G \in (0, 1)$ be a row player's belief that his opponent will play L and $\beta_c^G \in (0, 1)$ be a column player's belief that his opponent will play U . If $\beta_r^1 > \beta_r^3$ and $\beta_c^1 > \beta_c^3$, there exist utility functions v_j for all θ_j , which represent \succsim_j , such that if $\hat{a}_{\theta_j}^G = \{U, D\}$ and $\hat{a}_{\theta_j}^G = \{L, R\}$ respectively, then θ_j is indifferent between his pure*

actions in the Bayesian game that results from G .

Proof. At first, consider the types θ_1 to θ_3 . The action sets of the types θ_1 and θ_2 in game 1 are $\hat{a}_{\theta_j}^G = \{U, D\}$, $j = 1, 2$. Given a belief $\beta_r^1 \in (0, 1)$ in the Bayesian game of game 1, there exist utility functions, which represent \succsim_1 and \succsim_2 , such that the types $j = 1, 2$ are indifferent if

$$\beta_r^1 = \frac{v_j(7, 7) - v_j(3, 8)}{(v_j(4, 4) - v_j(8, 3) + v_j(7, 7) - v_j(3, 8))}. \quad (7)$$

In game 3, the action set of all three types is $\{U, D\}$. Hence, given a belief $\beta_r^3 \in (0, 1)$ in the Bayesian game 3, there exists utility functions for $j = 1, 2, 3$ such that the types are indifferent if

$$\beta_r^3 = \frac{v_j(7, 7) - v_j(3, 3)}{(v_j(4, 4) - v_j(8, 3) + v_j(7, 7) - v_j(3, 3))}. \quad (8)$$

Finally, in game 2, only θ_1 is a relevant type. The action set of θ_1 in game 2 is $\{U\}$. That means, given a belief $\beta_r^2 \in (0, 1)$, a utility function v_1 that represents \succsim_1 needs to satisfy:

$$\beta_r^2 v_1(5, 8) + (1 - \beta_r^2) v_1(7, 7) \geq \beta_r^2 v_1(6, 2) + (1 - \beta_r^2) v_1(3, 3) \quad (9)$$

Now, choose $v_j(7, 7)$, $v_j(8, 3)$, and $v_j(4, 4)$ such that the utility values are consistent with the ordering \succsim_j for $j = 1, 2, 3$. From equation (7) and (8), we obtain

$$v_j(3, 8) - v_j(3, 3) = \frac{(v_j(7, 7) - v_j(8, 3))(\beta_r^1 - 1)(\beta_r^3 - \beta_r^1)}{\beta_r^1 q_r^3}.$$

Since $\beta_r^1 > \beta_r^3$, and for all three types, $j = 1, 2, 3$, $v_j(7, 7) > v_j(8, 3)$, we have that $v_j(3, 8) > v_j(3, 3)$, which is consistent with the rankings of θ_1 , θ_2 and θ_3 given in

Table 9. Note that we have now shown that the lemma holds for the types θ_2 and θ_3 since there are no further restrictions concerning their utility functions. For type θ_1 , we may define a utility function v_1 , which represents \succsim_1 , such that the distance $v_1(6, 2) - v_1(5, 8)$ is arbitrary small. Hence, the lemma is also true for θ_1 .

Now, consider the types θ_4 to θ_6 . We omit the obvious proof for type θ_4 , and turn to the types θ_5 and θ_6 . In game 4, the action set of both types is $\{U, D\}$. Given a belief $\beta_r^4 \in (0, 1)$, the indifference condition for $j = 5, 6$ in the Bayesian game 4 is

$$\beta_r^4 = \frac{v_j(3, 8) - v_j(2, 2)}{(v_j(3, 8) - v_j(2, 2) + v_j(8, 3) - v_j(7, 7))}. \quad (10)$$

In game 2, the action set of both types is $\{U\}$. Hence, given $\beta_r^2 \in (0, 1)$, for $j = 5, 6$,

$$\beta_r^2 v_j(5, 8) + (1 - \beta_r^2) v_j(7, 7) \geq \beta_r^2 v_j(6, 2) + (1 - \beta_r^2) v_j(3, 3). \quad (11)$$

If there is a strict inequality in equation (11), it is obvious that one can choose utility values that satisfy (10) and (11) and represent \succsim_5 and \succsim_6 . Suppose that (11) holds with equality and choose the utility values $v_j(5, 8)$, $v_j(6, 2)$, $v_j(8, 3)$, $v_j(2, 2)$, and $v_j(3, 8)$ for $j = 5, 6$ such that the utility functions represent the ordinal rankings \succsim_5 and \succsim_6 . Then, equation (10) and (11) imply that $v_j(7, 7) > v_j(3, 3)$. This is consistent with \succsim_5 and \succsim_6 , which shows that the lemma holds for θ_5 and θ_6 .

For the column types θ_7 to θ_{11} , the lemma can be proved similarly to the proof of the row types θ_1 to θ_6 . □

For the Bayesian game of each game G , consider a strategy profile σ_G^* such that

- (i) $\sigma_G^*(\theta_j)$ is a proper mixed action for non-noisy row types θ_j where $\hat{a}_{\theta_j}^G = \{U, D\}$.
- (ii) $\sigma_G^*(\theta_j)$ is a proper mixed action for non-noisy column types θ_j where $\hat{a}_{\theta_j}^G =$

$\{L, R\}$.

(iii) $\beta[L | \sigma_1^*] > \beta[L | \sigma_3^*]$ and $\beta[U | \sigma_1^*] > \beta[U | \sigma_3^*]$.

It is obvious that such a strategy profiles exist. By assumption, noisy types randomly select a pure action, i.e., they play a proper mixed action. Given any mixed action of noisy types in the Bayesian game of each game G , define utility functions for all non-noisy types such that these types have no incentive to deviate from σ_G^* in each game. By Lemma 1, we know that such utility functions always exist. Then, σ_G^* is a QBNE for the Bayesian game of G for $G \in \{1, 2, 3, 4\}$. One can easily check that $\hat{a}^G \subseteq \text{supp}(\sigma_G^*)$ for all $G \in \{1, 2, 3, 4\}$, which proves the proposition. \square

A.2 Additional Data

A.2.1 Preferences, types, and associated action sets

Tables 9 and 10 show how subjects ranked the 8 payment pairs presented to them in stage 1 of the experiment. Payment pairs that are assigned a lower number are preferred to payment pairs with a higher number. There were 82 row and 82 column players and the rightmost column indicates the number of subjects who ordered the payment pairs in the corresponding way. The types that are relevant for the proof of Proposition 2 are labeled with names. Table 11 shows in which games the relevant types have no strictly dominant action and the observed action sets associated with these types per game.

Table 9: Preferences of row players, both treatments

Type name	(8,3)	(7,7)	(5,8)	(4,4)	(6,2)	(3,8)	(3,3)	(2,2)	n
θ_5	1	2	4	5	3	6	7	8	28
θ_6	1	2	4	5	3	7	6	8	10
θ_1	1	2	4	5	3	6	6	7	6
	2	1	4	5	3	6	7	8	6
θ_2	1	2	3	5	4	6	7	8	4
	2	1	3	5	4	6	7	8	4
θ_3	2	1	3	4	5	6	7	8	2
	1	2	3	4	3	5	6	7	2
θ_4	2	1	3	6	4	5	7	8	2
	1	2	3	4	5	6	7	8	2
	1	2	4	6	3	5	7	8	2
	2	1	4	5	3	6	6	7	1
	2	1	5	3	4	6	7	8	1
	1	3	4	4	2	6	5	7	1
	1	1	2	3	2	3	4	5	1
	3	1	2	4	6	5	7	8	1
	3	1	2	4	4	3	5	6	1
	5	1	7	2	6	8	3	4	1
1	2	5	4	3	7	6	8	1	
6	1	2	3	5	6	7	8	1	
1	2	4	5	3	7	7	8	1	
1	2	7	4	3	8	5	6	1	
3	1	2	5	5	3	7	8	1	
1	1	3	4	2	5	6	7	1	
1	2	4	5	3	6	6	8	1	

Table 10: Preferences of column players, both treatments

Type name	(8,3)	(7,7)	(8,5)	(4,4)	(2,6)	(3,8)	(3,3)	(2,2)	n
θ_9	2	3	1	4	7	5	6	8	35
θ_7	3	2	1	4	7	5	6	8	6
	3	2	1	5	7	4	6	8	4
	3	2	1	5	6	4	7	8	3
θ_8	3	1	2	4	7	5	6	8	3
	3	1	2	4	8	6	5	7	2
θ_{11}	1	2	1	3	6	4	5	7	2
θ_{10}	1	3	2	4	8	6	5	7	2
	2	3	1	4	7	6	5	7	1
	3	1	2	5	6	4	7	8	1
	4	1	3	2	8	6	5	7	1
	1	3	2	4	6	5	7	8	1
	3	1	2	4	6	5	7	8	1
	4	1	2	3	8	6	5	7	1
	2	1	1	3	5	4	4	5	1
	2	3	1	4	6	5	7	8	1
	2	3	1	4	8	7	5	6	1
	1	2	1	3	5	4	4	6	1
	5	1	4	2	8	6	3	7	1
	2	2	1	3	6	4	5	7	1
	3	1	2	4	8	7	5	6	1
	2	1	1	3	5	2	4	6	1
	2	1	3	4	6	5	7	8	1
	1	2	1	3	5	4	4	5	1
	3	1	2	4	6	5	5	6	1
	4	1	2	3	5	4	6	7	1
	1	2	1	3	7	5	4	6	1
	3	2	1	4	8	6	5	7	1
	2	3	1	4	6	5	5	6	1
	1	3	4	5	5	7	7	8	1
	1	3	2	4	8	7	5	6	1
	4	1	3	2	7	5	6	8	1
	2	3	4	5	6	6	7	1	1

Table 11: Types and associated action sets per game

Type name	#Subjects	No strictly dominant action in game	Observed action sets
θ_1	6	G=1,2,3	$\hat{a}_{\theta_1}^1 = \{U, D\}, \hat{a}_{\theta_1}^2 = \{U\}, \hat{a}_{\theta_1}^3 = \{U, D\}$
θ_2	4	G=1,3	$\hat{a}_{\theta_2}^1 = \{U, D\}, \hat{a}_{\theta_2}^3 = \{U, D\}$
θ_3	2	G=1,3	$\hat{a}_{\theta_3}^1 = \{D\}, \hat{a}_{\theta_3}^3 = \{U, D\}$
θ_4	2	G=1,2,4	$\hat{a}_{\theta_4}^1 = \{U\}, \hat{a}_{\theta_4}^2 = \{U\}, \hat{a}_{\theta_4}^4 = \{D\}$
θ_5	28	G=2,4	$\hat{a}_{\theta_5}^2 = \{U\}, \hat{a}_{\theta_5}^4 = \{U, D\},$
θ_6	10	G=2,4	$\hat{a}_{\theta_6}^2 = \{U\}, \hat{a}_{\theta_6}^4 = \{U, D\},$
θ_7	6	G=1,2,3	$\hat{a}_{\theta_7}^1 = \{L, R\}, \hat{a}_{\theta_7}^2 = \{L, R\}, \hat{a}_{\theta_7}^3 = \{L, R\},$
θ_8	3	G=1,3	$\hat{a}_{\theta_8}^1 = \{L, R\}, \hat{a}_{\theta_8}^3 = \{L, R\},$
θ_9	35	G=2,3,4	$\hat{a}_{\theta_9}^2 = \{L, R\}, \hat{a}_{\theta_9}^3 = \{L, R\}, \hat{a}_{\theta_9}^4 = \{L, R\},$
θ_{10}	2	G=2,3	$\hat{a}_{\theta_{10}}^2 = \{L, R\}, \hat{a}_{\theta_{10}}^3 = \{L, R\}$
θ_{11}	2	G=3,4	$\hat{a}_{\theta_{11}}^3 = \{R\}, \hat{a}_{\theta_{11}}^4 = \{L\}$

A.2.2 Details of the robustness check tests for the main result

Tables 12 and 13 report the results of a two-tailed Fisher exact test of the null hypothesis that the probability that a subject plays the equilibrium strategy is the same in both treatments. These tests were run separately for each of the 4 games. n_base is the number of observations in treatment baseline and n_info the number of observations in treatment info. The tests reported in Table 12 include all subjects while those reported in Table 13 include consistent subjects only.

Table 12: Fisher exact test (two-tailed), all subjects.

Game	n_base	n_info	p-value
1	19	12	0.452
2	11	17	1.000
3	27	28	0.014
4	18	15	0.296

Table 13: Fisher exact test (two-tailed), consistent subjects only.

Game	n_base	n_info	p-value
1	16	10	0.109
2	7	13	0.374
3	21	22	0.045
4	10	9	0.170

A.3 Equilibrium under ambiguity

The concept of Eichberger and Kelsey (2014) is called “equilibrium under ambiguity (EUA)”. In their concept, player i ’s beliefs about the behavior of other players is represented by a capacity ν_i defined on $S_{-i} = \times_{j \in I \setminus \{i\}} S_j$, where S_j is the set of player j ’s pure strategies. Given his beliefs ν_i , player i ’s payoff from a pure strategy $s_i \in S_i$ corresponds to the Choquet integral of his payoff function $u_i(s_i, s_{-i})$ with respect to ν_i :

$$\begin{aligned} V_i(s_i, \nu_i) &= \int_{S_{-i}} u_i(s_i, s_{-i}) d\nu_i \\ &= u_i(s_i, s_{-i}^1) \nu(s_{-i}^1) + \sum_{r=2}^R u_i(s_i, s_{-i}^r) [\nu(s_{-i}^1, \dots, s_{-i}^r) - \nu(s_{-i}^1, \dots, s_{-i}^{r-1})], \end{aligned}$$

where the strategy combinations in S_{-i} are numbered so that $u_i(s_i, s_{-i}^1) \geq u_i(s_i, s_{-i}^2) \geq \dots \geq u_i(s_i, s_{-i}^R)$. Player i ’s best responses to his belief ν_i are defined in the usual way as

$$R_i(\nu_i) = \{s_i \mid s_i \in \arg \max_{s_i \in S_i} V_i(s_i, \nu_i)\}.$$

An essential ingredient of the model is the notion of support for a non-additive measure. Eichberger and Kelsey define the support of a convex capacity as the inter-

section of the supports of the probability measures in the core of the capacity:¹⁰

Definition 10. *The support of a convex capacity μ on S_{-i} is defined as*

$$\text{supp}(\mu) = \bigcap_{\pi \in \text{core}(\mu)} \text{supp}(\pi).$$

It is well-known that convex capacities represent ambiguity-aversion. To capture optimistic behavior, Eichberger and Kelsey use the class of capacities introduced by Jaffray and Philippe (1997) (JP-capacities). A JP-capacity has convex and concave parts. It is defined as a mixture of a convex capacity with its dual capacity.¹¹ Eichberger and Kelsey define the support of a JP-capacity ν , $\text{supp}_{JP}(\nu)$, as the support of its convex part according to Definition 10. This support definition has a useful implication for neo-additive capacities introduced by Chateauneuf et al. (2007):

Proposition 3 (Eichberger and Kelsey, 2014). *Let $\nu = \delta\alpha + (1-\delta)\pi$ be a neo-additive capacity on S_{-i} , where $\alpha, \delta \in [0, 1]$, then $\text{supp}_{JP}(\nu) = \text{supp}(\pi)$.*

We will use neo-additive capacities to discuss the example in this section.

An equilibrium under ambiguity is a belief system in which, for each player i , the nonempty support of player i 's belief about the opponents' behavior lies in the Cartesian product of the opponents best responses given their beliefs about the behavior of other players. To put it differently, in an equilibrium under ambiguity, the beliefs that agents hold are reasonable in the sense that neither player expects other players to play strategies that are not best responses given their beliefs.

Definition 11. A belief system (ν_i^*, ν_{-i}^*) is *an equilibrium under ambiguity* if for all $i \in I$

¹⁰For alternative support definitions and for arguments supporting Definition 10, see Eichberger and Kelsey (2014).

¹¹The dual capacity of capacity μ is defined by $\bar{\mu}(E) = 1 - \mu(E^c)$. Hence, if μ is convex, then $\bar{\mu}$ is concave.

$$\text{supp}(\nu_i^*) \subseteq \bigtimes_{j \in I \setminus \{i\}} R_j(\nu_j^*) \text{ and } \text{supp}(\nu_i^*) \neq \emptyset.$$

In what follows, we show that an equilibrium under ambiguity can explain maximin and maxmax strategy choices.

Example 2. *Consider game 3 in stage 2 of treatment info. Suppose that the game is played by two subjects whose utility functions correspond to their own payments. That is, the game takes the following form:*

	<i>L</i>	<i>R</i>
<i>U</i>	4, 4	8, 3
<i>D</i>	3, 3	7, 7

Obviously, the row player has a strictly dominant strategy (*U*). If the column player believes that row will pick *U*, he will play *L*. The game has a unique Nash equilibrium (*U, L*). Whether or not the Nash equilibrium is played depends on the belief of the column player about whether the row player behaves rationally, i.e., whether row will play the strictly dominant strategy.

Denote the players by $I = \{r, c\}$, where *r* stands for row and *c* for column. If the column player is not sure whether row behaves rationally, he may try to reach the highest possible outcome (7) by playing strategy *R*. To show that this strategy choice is consistent with an equilibrium under ambiguity, suppose that the beliefs of the column player about row's behavior can be represented by a neo-additive capacity ν_c with reference prior $\pi_c = (\pi_c(U), \pi_c(D)) = (1, 0)$. This can be viewed as a situation where column is uncertain about the prior π_c , i.e., whether row plays *U*. Furthermore, let column be an ambiguity-loving player, for simplicity, assume that $\alpha_c = 0$. We may interpret the parameter δ_c as the degree of ambiguity about π_c . The higher δ_c , the higher the degree of ambiguity. Given this belief, column's payoff from *L* equals

$$V_c(L, \nu_c) = (1 - \delta_c) \cdot 4 + \delta_c \cdot (\max\{u_c(L, s_r) \mid s_r \in S_r\}) = 4,$$

and column's payoff from R is

$$V_c(R, \nu_c) = (1 - \delta_c) \cdot 3 + \delta_c \cdot (\max\{u_c(R, s_r) \mid s_r \in S_r\}) = 3 + 4\delta_c.$$

Hence, if column is sufficiently uncertain about π_c ($\delta_c > \frac{1}{4}$), he will choose strategy R .

Suppose that the beliefs of the row player about column's behavior can also be represented by a neo-additive capacity ν_r with reference prior $\pi_r = (\pi_r(L), \pi_r(R)) = (0, 1)$. It is straightforward that the row player will play U given such a belief. Taken together, for $\delta_c > \frac{1}{4}$, we have that

$$R_r(\nu_r) = U \text{ and } R_c(\nu_c) = R,$$

and, by Proposition 3 it holds that

$$\text{supp}_{JP}(\nu_r) = \text{supp}(\pi_r) = R \text{ and } \text{supp}_{JP}(\nu_c) = \text{supp}(\pi_c) = U.$$

Consequently, the system (ν_r, ν_c) is an equilibrium under ambiguity in which the column player plays the maxmax strategy R . Similarly, one can show that the equilibrium under uncertainty concept can rationalize maxmin behavior if the players are ambiguity-averse.

References

- AUMANN, R. AND A. BRANDENBURGER (1995): “Epistemic conditions for Nash equilibrium,” *Econometrica*, 5, 1161–1180.
- CHATEAUNEUF, A., J. EICHBERGER, AND S. GRANT (2007): “Choice under uncertainty with the best and worse in mind: Neo-additive capacities,” *Journal of Economic Theory*, 137, 538–567.
- DUFWENBERG, M. AND G. KIRCHSTEIGER (2004): “A theory of sequential reciprocity,” *Games and Economic Behavior*, 47, 268–298.
- EICHBERGER, J. AND D. KELSEY (2000): “Non-Additive beliefs and strategic equilibria,” *Games and Economic Behavior*, 30, 183–215.
- (2014): “Optimism and pessimism in games,” *International Economic Review*, 55, 483–505.
- ELLSBERG, D. (1961): “Risk, ambiguity, and the Savage axioms,” *Quarterly Journal of Economics*, 75, 643–669.
- FISCHBACHER, U. (2007): “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, 10, 171–178.
- GREINER, B. (2004): “An online recruitment system for economic experiments,” in *Forschung und Wissenschaftliches Rechnen 2003. Gesellschaft für wissenschaftliche Datenverarbeitung Bericht, Vol. 63*, ed. by K. Kremer and V. Macho, Göttingen, Germany: Gesellschaft für wissenschaftliche Datenverarbeitung, 79–93.
- HARSANYI, J. (1967-1968): “Games with incomplete information played by ‘Bayesian’ players, Parts I-III.” *Management Science*, 14, 159–182, 320–334, 486–502.
- HEALY, P. (2011): “Epistemic foundations for the failure of Nash equilibrium,” Mimeo.
- HU, A. AND L. ZOU (2015): “Sequential auctions, price trends, and risk preferences,” *Journal of Economic Theory*, 158, 319–335.
- JAFFRAY, J. AND F. PHILIPPE (1997): “On the existence of subjective upper and lower probabilities,” *Mathematics of Operations Research*, 22, 165–185.
- LEHRER, E. (2012): “Partially specified probabilities: Decisions and games,” *American Economic Journal: Microeconomics*, 4, 70–100.
- LO, K. (1996): “Equilibrium in beliefs under uncertainty,” *Journal of Economic Theory*, 71, 443–484.

- NASH, J. (1950): “Equilibrium points in n-person games,” *Proceedings of the National*, 36, 48–49.
- (1951): “Non-cooperative games,” *The Annals of Mathematics*, 54, 286–295.
- RABIN, M. (1993): “Incorporating fairness into game theory and economics,” *American Economic Review*, 83, 1281–1302.
- WOLFF, I. (2014): “When best-replies are not in equilibrium: Understanding cooperative behaviour,” Mimeo.