Identifying the reasons for coordination failure in a laboratory experiment*

Philipp Külpmann[†] Davit Khantadze[‡]

November 8, 2018

Abstract

We investigate the effect of absence of common knowledge on the outcomes of coordination games in a laboratory experiment. Using cognitive types, we can explain coordination failure in pure coordination games while differentiating between coordination failure due to first- and higher-order beliefs.
In our experiment, around 76% of the subjects have chosen the payoff-dominant equilibrium strategy despite the absence of common knowledge. However, 9% of the players had first-order beliefs that lead to coordination failure and another 9% exhibited coordination failure due to higher-order beliefs. Furthermore, we compare our results with predictions of commonly used models of higher-order beliefs.

JEL codes: C72, C92, D83
Keywords: Higher-order beliefs, coordination failure, cognitive abilities, experimental economics, game theory

*We would like to thank Andreas Blume, Zahra Gambarova, Wieland Müller, Mahnaz Nazneen, Eugenio Proto, Andis Sofianos and especially Christoph Kuzmics and Daniel Sgroi for valuable comments and suggestions. Philipp acknowledges support through grants from the German Research Foundation (DFG, grant number: KU 3071/1) and the German Academic Exchange Service (DAAD). Furthermore, he would like to thank the Department of Economics of the University of Warwick, where this research was carried out, for its hospitality.

[†]University of Vienna, philipp.kuelpmann@univie.ac.at
[‡]University of Warwick, d.khantadze@warwick.ac.uk

1. Introduction

If you have lost your spouse in a department store and both of you are trying to find each other, the answer to the (seemingly simple) question of "Will she look for me at the coffee bar or at the exit?" depends not only on the answer to the question "Does she think I am looking for her at the coffee bar or at the exit?" (i.e., something we will call the first-order belief) but also on the answers to "Does she think that I think that she thinks that I am looking for her at the coffee bar or at the exit?" (i.e., the second-order belief or "What is her first-order belief?") and on infinitely more levels of beliefs. This paper addresses the question if people actually use beliefs of a higher order.

When modeling human behavior, most works assume that players have common knowledge about the structure of the game, i.e., that all players know the structure, that all players know that everyone else knows the structure and so on.[1] Furthermore, we assume that players do not only have common knowledge about publicly known properties of the game but also about the distributions of unknown factors of the game, like the other players' types (for example if I'd rather wait at the coffee bar or the exit). The absence of common knowledge leads to complex belief hierarchies, so called *higher-order beliefs*. The first level of these beliefs, so called first-order beliefs, might be a belief over the other player's type. A second-order belief would then be a belief over the belief of the other player about your type (i.e., a belief over the other player's first-order belief) and so on ad infinitum.

In the game theoretical literature many different assumptions and models of higher-order beliefs exist and many of these lead to very different predictions even in simple games like the pure coordination game we are using in this paper.[2] The question, what kind of model of higher-order beliefs players actually use, seems to be an empirical question which we are trying to address in this paper.

To do so, we take up the experimental results and setup of Blume and Gneezy (2010), in which there is an issue of cognitive difficulties, to analyze the effects of higher-order beliefs. Blume and Gneezy (2010) used a slightly difficult coordination game (the so-called 5-sector disc), in which there is a better (i.e., risk- and payoff-dominant) option which is harder to find. They were able to show that participants form beliefs about the cognitive abilities of other participants and, if these beliefs are pessimistic, they hinder coordination between the players (i.e.,

[1] However, there is a (mostly theoretical) literature on universal type spaces, introduced by Harsanyi (1967/68) and formalized by Mertens and Zamir (1985) and Brandenburger and Dekel (1993), analysis games without common knowledge.

[2] A brief overview of some models of higher-order beliefs can be found in Section 7 and a more detailed discussion in Section 5.

that "beliefs matter"). However, they have not taken into account the effect of higher-order beliefs about cognitive abilities. Therefore, we modify their experimental setup in order to distinguish the effect of first-order beliefs players form about the cognitive ability of their opponents (i.e., if players trust in the cognitive ability of their partners) and higher-order beliefs.

We introduce a new treatment in which participants guess what other participants play against themselves. This allows us to identify first-order beliefs and therefore separate first- from higher-order beliefs.

Using the data from these treatments, we can answer the following three questions:

- **Are players able to coordinate in the absence of common knowledge?**
- Can coordination fail because players underestimate the skill of the other players? Or, in other words, **do first-order beliefs matter?**
- Can coordination fail because players think "too much" about what others might think? Or, in other words, **do higher-order beliefs matter?**

Using Blume and Gneezy's (2000) 5-sector disc, we were able to find answers to all three questions: In the experiment, we were able to reproduce Blume and Gneezy's (2010) result, that the majority of players had no problem choosing the Pareto-dominant equilibrium strategy of the game (i.e., coordination is possible). Furthermore, some players switch to the worse equilibrium strategy because of first- and higher-order beliefs (i.e. first- and higher-order beliefs matter).

More important applications than the search for one's husband or wife in a department store are suggested by recent studies in sociology and development studies, like Bicchieri (2005). She claims that common knowledge plays a significant role in the fight against female genital mutilation. Our results might help to improve our understanding of why some organizations are significantly more successful in the fight against female genital mutilation (FGM) than others. In Section 6, we discuss this application in more detail and compare it to the fight against foot binding.

Apart from the FGM application, higher-order beliefs have applications in many different fields. For example, they might be a reason for bank runs (i.e., the belief that a bank will be fine but a fear that others might think that the bank is in trouble might cause a bank run on this bank), arms races and financial crises.

2. Experimental design

Measuring higher-order beliefs is very complicated, as there is an "uncertainty principle" (as already discussed by Blume and Gneezy (2010)) at work; i.e., it is

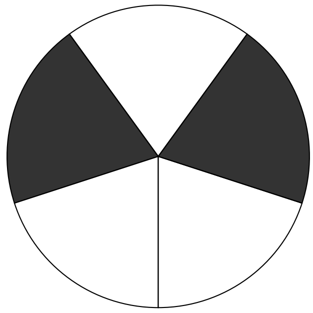

Figure 1: 5-sector-disc

hard to measure beliefs without introducing or changing them.[3] Furthermore, introducing absence of common knowledge is difficult; When told that they are given a random number, subjects usually assume that it is drawn from a uniform distribution. Explicitly stating that the distribution is unknown leads to a myriad of other problems. Subjects could, for example, assume a strategic selection of the distribution by the experimenter. Finally, we need to have some sort of control over the fraction of high-cognition players, so that the action only available to the high-cognition players is the one with the highest expected payoff (see Section 5).

We solve all three problems by utilizing Blume and Gneezy's (2000) 5-sector disc. This is a disc with 5 equally large sectors on it, 2 black and 3 white, as depicted in Figure 1.[4]

The disc has the same sectors on the front- and backside of the disc and can be flipped and rotated.

As the disc can be flipped, the subjects face symmetry constraints and can therefore not distinguish all five sectors. These symmetries cannot be overcome and therefore not all Nash equilibria are possible given the particular frame. Only certain "attainable" equilibria are possible, as defined originally in Crawford and Haller (1990), and further developed by Blume (2000) and Alós-Ferrer and Kuzmics (2013).

The property of this disc which is most important for this paper is that it has a single distinct white sector: The sector adjacent to both black sectors (Figure 1).[5]

For the subjects there are then, in principle, three distinguishable sets of sectors: the black sectors (B), the uniquely identifiable white sector (D), and the other white

[3] Either by making the subjects realize that there might be something like a higher-order belief or by them trying to be a good subject (Orne (1962)). A more extensive discussion of this uncertainty principle can be found in Appendix B.

[4] There is a second version of this disc, with a significantly harder to find distinct sector, with adjacent black sectors. However, for this disc, the fraction of players who were able to identify the distinct sector is too small (i.e., not satisfying the conditions derived in Section 5).

[5] More about the properties of this disc can be found in Blume and Gneezy (2000).

sectors (W').

The key assumption behind the experiment (and also behind Blume and Gneezy (2000) and Blume and Gneezy (2010) and very much supported by their findings), is that not all subjects realize that there is a uniquely identifiable sector, which leads to two different cognitive types, the high type, who can identify the distinct sector, and the low type, who cannot.

The low type then faces an additional symmetry constraint and has only two distinguishable sets of sectors to choose from: One of the two black sectors (B) or one of the three white sectors (W).

Note that the lower type has no knowledge about the existence of another type or the distinct sector.

The subjects then played three treatments in a random order without feedback after hearing and reading the instructions and completing an extensive quiz:[6]
The **Self Treatment** in which the subject gets the disc twice, every time randomly turned and rotated, and gets £5 if she picks the same sector twice.
In the **Prediction Treatment** one subject (she) is told that another subject (he) plays the *Self Treatment* (with a possibly differently turned and rotated disc). She has to pick one sector and every time he picks the sector she picked, she gets £2.5.[7]
Finally, the **Coordination Treatment**, in which two players pick simultaneously a sector on a (randomly turned and rotated) disc and, if both players pick the same sector, both receive £5.

2.1. Predictions

How can we use this design to test the three initial questions stated in the introduction? Let us have a look how we expect low- and high-cognition players to behave in the three different treatments.

In the Self Treatment a high-cognition player has 9 possible choices: She can pick any of three actions (D, B, W') in the first stage and then pick any of the three actions in the second stage. This decision problem for the high-cognition player has a unique optimal solution: pick the distinct sector twice, giving her a probability to win of 1.

A low-cognition player is only aware of four possible choices: He can pick B or W in the first stage and then pick B or W in the second stage. The low-cognition

[6] For the complete instructions and a description of the quiz see the Online Appendix: `http://www.kuelpmann.org/research/`.

[7] Adding another treatment in which subjects have to predict what another subjects does in the Prediction Treatment would, in theory, allow to explicitly check for second-order beliefs (or, when repeating this any higher-order belief). However, don't belief this will work with the 5-sector disc, as it probably requires too much attention and mental effort which most subjects might not be willing to exert.

player also has a unique optimal choice: pick B in both stages, giving him a probability to win of $\frac{1}{2}$.

Therefore, we would expect a high-cognition player to choose the distinct sector twice and a low-cognition player to pick a black sector twice.

In the Prediction Treatment, the action taken by a subject should only depend on her type and her first-order belief about the type of the other player. A low-cognition player will always choose B, whereas a risk-neutral, high-cognition player should pick D if his belief that the other player is also of the high type is at least $\frac{1}{3}$ and B otherwise.[8]

The coordination treatment is best depicted as a bi-matrix game with three (for the high-cognition player) and two (for the low-cognition player) pure strategies, with winning probabilities as depicted in Figure 2 and Figure 3. We expect a low-cognition player to play B, as it is the payoff- and risk-dominant equilibrium, whereas a high-cognition player's choice depends on her belief hierarchy: If anywhere in her complete hierarchy a belief lower than $\frac{1}{3}$ (or $\frac{1}{2}$ for very risk averse players) that the other player is a high-cognition player or that the other player thinks that she is a high-cognition player, ... (or, in short, that there is no common-p belief among the high-cognition players of $\frac{1}{3}$ or higher, that both players are high-cognition players), she will choose B, otherwise she will choose D.

	W'	B	D
W'	$\frac{1}{2}$	0	0
B	0	$\frac{1}{2}$	0
D	0	0	1

Figure 2: High-cognition player winning probabilities

	W	B
W	$\frac{1}{3}$	0
B	0	$\frac{1}{2}$

Figure 3: Low-cognition player winning probabilities

Unfortunately, neither the theoretical nor the experimental literature on higher-order beliefs can tell us which of the two will be chosen. Even small variations in the assumptions of theoretical models of higher-order beliefs can generate both equilibria. Therefore, this question seems to be an empirical one, which we are trying to answer in this paper. However, a more detailed explanation of how different models of higher-order beliefs work in our game can be found in Section 5.

[8]Allowing for risk-averse players, this fraction has to be between $\frac{1}{3}$ and $\frac{1}{2}$, depending on the degree of risk aversion.

2.2. Hypotheses

Using our design, we can formulate three hypotheses to test the three research questions stated earlier. In the following we will use a shorthand for players' strategies such as: "W'W' B D" means that a player selected one of the two white sectors twice in the Self Treatment, one of the black sectors in the Prediction Treatment and the distinct sector in the Coordination Treatment.

The answer to our first question "Is coordination possible? " or, in the words of our model "Do high-cognition players use the first-best strategy a_m despite the absence of common knowledge? " is suggested by the literature on focal points (e.g., Sugden (1995) or Crawford, Gneezy, and Rottenstreich (2008)) and supported by the experimental literature on coordination games (e.g., Van Huyck, Battalio, and Beil (1990) or Cooper, DeJong, Forsythe, and Ross (1990)):

Hypothesis 1 (Coordination is possible)**.** *High-cognition players choose in the Coordination Treatment D more often than any other choice.*

We are using a within-subject design to test the hypotheses: Only high-cognition players can identify the best equilibrium, so we don't have to consider other types. We can identify these players with the help of the the Self Treatment. If high-cognition players, i.e., the ones who have been able to identify "D" in the Self Treatment, coordinate on *D* in the Coordination Treatment we know that coordination is possible, even in the absence of common knowledge.

The next two hypotheses extend on Blume and Gneezy's (2010) hypothesis that "beliefs matter": Hypothesis 2 formalizes the question "Does coordination fail because some high-cognition players underestimate the fraction of high-cognition players? " or "Is there coordination failure due to first-order beliefs? "

Hypothesis 2 (First-order beliefs matter)**.** *There are high-cognition subjects who choose a black sector in the Prediction Treatment and Cooperation Treatment, i.e., play "DD B B".*

We already know that we can identify players' types with the help of the Self Treatment. Furthermore, the Prediction Treatment identifies players who think that more than $\frac{1}{3}$ of the other players can not identify the distinct sector.

Most of the problems in models of higher-order beliefs stem from the fact that there are infinitely many levels of beliefs. However, evidence from the laboratory indicates that people are not able to use higher-order rationality,[9] a requirement for coordination problems due to higher-order beliefs. Furthermore, even in studies of level-k reasoning, where players are framed and incentivized on using higher-order beliefs, players still rarely use high levels of reasoning.[10]

[9] Kneeland (2015) shows that only about 22% of all players use more than third-order rationality.

[10] In Arad and Rubinstein's (2012) 11-20 game, 80 % of the players only use 3rd-order beliefs or lower despite the game being designed to facilitate higher-order reasoning.

Therefore, the third question, if there is is coordination failure due to higher-order beliefs, or if high-cognition players use the first-best strategy a_m despite the absence of common knowledge, arises naturally:

Hypothesis 3 (Higher-order beliefs matter)**.** *There are high-cognition subjects who play the distinct sector in the Prediction Treatment and a black sector in the Cooperation Treatment, i.e., play "DD D B".*

Our design allows for another robustness check: There is an attainable strategy which is very similar to the one we use to identify first- and higher-order beliefs: "DD B D". This strategy will only be chosen if players belief that their partner is of the low type, but still plays "D" in the in the Coordination Treatment. This strategy can therefore not be explained using our model.

Hypothesis (Robustness check)**.** *"DD B D" is played less often than "DD B B" and "DD D B".*

3. The Model

Players face following coordination problem: if they pick the same sector from a five-sector disc they get a reward. Five-sector disc has three white and two black sectors. Players face symmetry constraints, which means that two black sectors are indistinguishable; three white sectors can be divided in the following groups: there is a distinct white sector and two indistinguishable white sectors. Call a player high-cognitive type if it is able to recognize a distinct sector and low-cognitive type otherwise.

Low and high cognitive players face different coordination problems: low-cognitive type is unaware of distinct sector and of high-cognitive types and faces following coordination game: coordinate on black or white sectors, with winning probabilities $\frac{1}{2}$ and $\frac{1}{3}$ respectively. High-cognitive player solves a coordination problem by being aware that it can face either high or low cognitive type.

Players get a monetary reward if coordination is successful. We assume that players have preferences that are represented by von Neumann/Morgenstern utility function and they value the monetary reward. It is without loss of generality to normalize the monetary reward to 1.[11] This implies that we can disregard players attitudes toward the risk and write player's payoff in terms of expected values.

[11]Consider von Neumann/Morgensern utility function u. Let's denote by s successful coordination, and by f - failure to coordinate. Then we can consider following transformation: $u'(s) = Au(s) + B = 1$ and $u'(f) = Au(f) + B = 0$; which implies solving two equations in two unknowns. The only assumption we make is that coordination payoff is preferred to the payoff when coordination fails.

First let us consider low-cognition players, who face following coordination game:

	W	B
W	$\frac{1}{3}$	0
B	0	$\frac{1}{2}$

This game has two Nash equilibria in pure strategies and one Nash equilibrium in mixed strategies. Note that (B, B) is payoff and risk dominant equilibrium. Equilibrium (B, B) is focal for low-cognition players and thus is more likely to be played. To simplify exposition we model low-cognition players as automaton choosing B with certainty. Our assumption is corroborated by the experiments.

Now we analyze the coordination game as faced by the high-cognition player.

Denote by p proportion of high-cognition players. With probability p high-cognition player is matched with another high-cognition player. In this case high-cognition player faces following coordination problem: if it beliefs it's partner chooses the distinct sector, then coordination is successful with probability 1 if it also chooses the distinct sector. Otherwise, if players coordinate on white or black sectors, probability of successful coordination is $\frac{1}{2}$.

With probability $(1 - p)$ high-cognition player is matched with low-cognition player. Because by assumption low-cognition player chooses black sector, this implies that high-cognition player can successfully coordinate only by choosing a black sector.

The following matrix depicts the game as faced by a high-cognition player - player 1 - where the two matrices give it's payoffs.

Coordination Game - High Cognition player

p

	W′	B	D
W′	$\frac{1}{2}$	0	0
B	0	$\frac{1}{2}$	0
D	0	0	1

1-p

	B
W′	0
B	$\frac{1}{2}$
D	0

If p is common knolwedge among high-cognitive players, then the set of Bayesian Nash equilibria for this game is a function of p. For $p < \frac{1}{3}$, B is a strictly dominant strategy for a high-cognition player and the unique Nash equilibrium of the game is (B, B).

For $\frac{1}{3} \leq p < \frac{1}{2}$, high-cognition player's game has two Nash equilibria in pure strategies: (B, B) and (D, D).

For $p \geq \frac{1}{2}$ the set of Nash equilibria is (W', W'), (B, B) and (D, D).

4. Results

The experiment was conducted at the DR@W Laboratory at the University of Warwick using the experimental software "z-Tree" developed by Fischbacher (2007). 130 subjects where recruited and received payments between £3 and £18. Before showing the results, let us briefly discuss the preliminaries of the experiment design.

The first preliminary is the focality of the distinct and the two black sectors. From the choice data in Figure 4 we can see that more than 95% of all players have chosen one of these sectors in the Coordination Treatment. Therefore, the black and distinct sectors seem to be focal in our game. The second preliminary is that there are enough high-cognition players, so that playing the high-cognition exclusive action is a payoff-dominant equilibrium for the players. In Figure 5 we can see that 58% of all players have chosen the distinct sector and are therefore considered high-cognition players. Therefore, playing the distinct sector would maximize the expected utility of high-cognition players in a game where the type distribution is common knowledge among high types independently of the degree of risk aversion (see Section 5). We can also see that the second most frequently observed behavior is choosing a black sector twice, whereas choosing a white sector twice (which includes choosing the distinct white sector once and another white sector once) and picking one black and one white sector (labeled “Other“) was very rare.

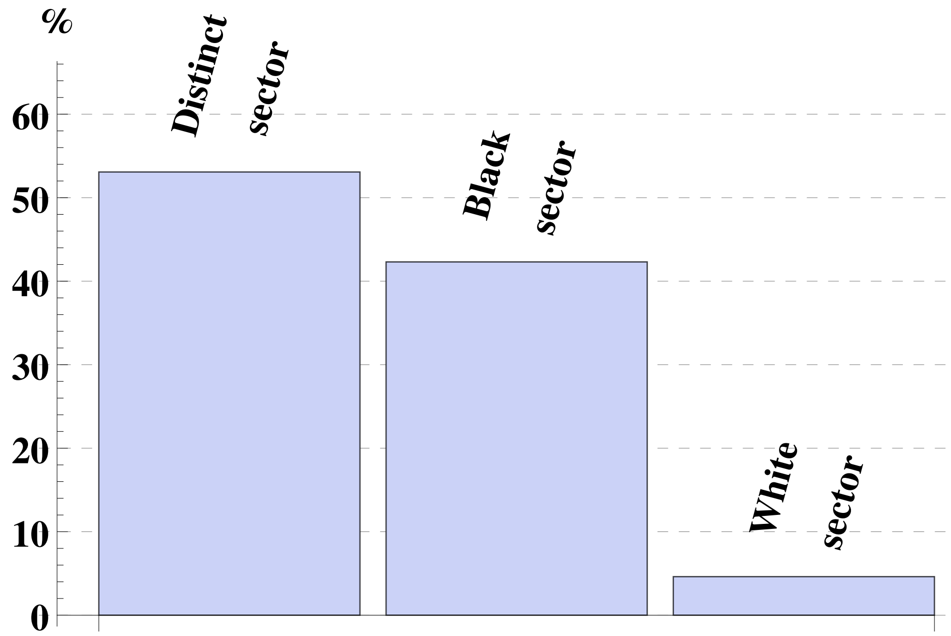

Figure 4: Results of the Coordination Treatment

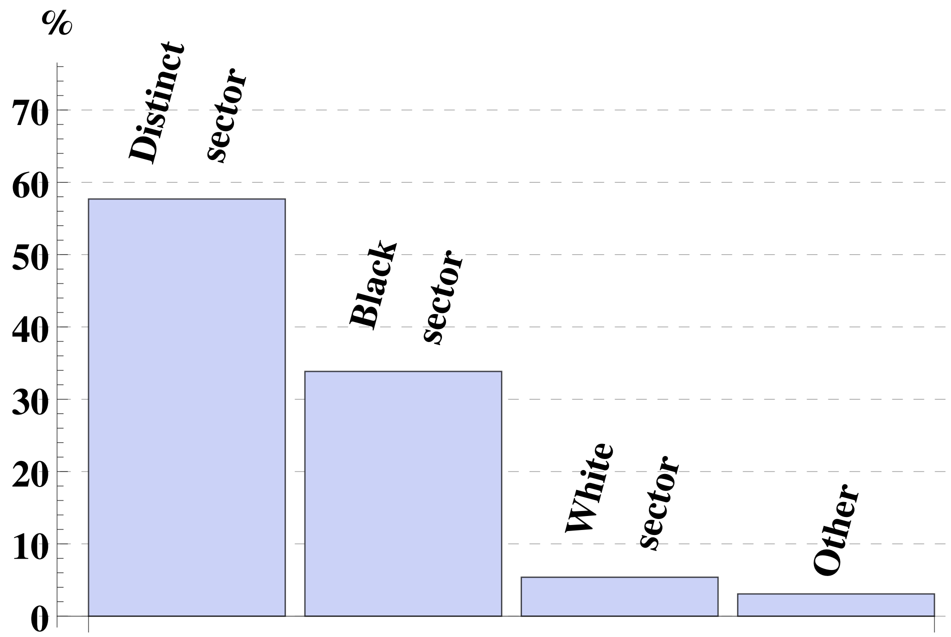

Figure 5: Results of the Self Treatment

These results are in line with Blume and Gneezy's (2010) results where around 52% (58% in our experiment) have been able to identify the distinct sector and around 23% (34%) have chosen the black sector. We contribute the significantly lower level of noise (8% vs 25%) to the extensive instructions and the quiz we conducted before the experiment.[12]

Due to the lower level of noise we are, unlike Blume and Gneezy (2010), able to use a within-subject design, in which each player has access to 625 possible strategies.[13] Of these strategies we consider 96.32% as "noise".[14] As the number of strategies which support our hypothesis are very low (1, 4 and 2 out of 625), the probability that someone chooses them by mistake is very low. For a detailed overview of all possible strategies and how we categorize them see Table 1.

[12]For the instructions and an overview of the quiz see the Online Appendix (`http://www.kuelpmann.org/research/`).

[13]We are here ignoring the order in which treatments are played.

[14]This noise includes not only players not understanding the experiment or behaving randomly but also "Eureka"-learning (which was a big problem in Blume and Gneezy (2010), see Appendix B), making a mistake (e.g., picking a not distinct white sector instead of the distinct sector, a mistake, which both of the authors made multiple times while testing the experiment) and beliefs of low-cognition players.

Description	Hypothesis	# of strategies	Proportion
DD D D	1: Coordination is possible	1	0.16%
DD B B	2: First-order beliefs matter	4	0.64%
DD D B	3: Higher-order beliefs matter	2	0.32%
BB B B	(Low-cognition players)	16	2.56%
"Noise"	-	602	96.32%
WW W W	(part of "Noise")	80	12.80%

Table 1: Overview of the strategies

Given the preliminaries, we can test hypotheses 1 through 3.

Hypothesis 1 (Coordination is possible)**.** *High-cognition players choose in the coordination treatment D more often than any other choice.*

The choice data from our experiment confirms this hypothesis. In Figure 6 we can see that 80% have chosen the strategy "DD D D". As this strategy represents only 0.16% of all available strategies (or 4% when excluding the Self Treatment), we can reject the null hypothesis of this high level of coordination being a result of random play ($p < 0.00001$).

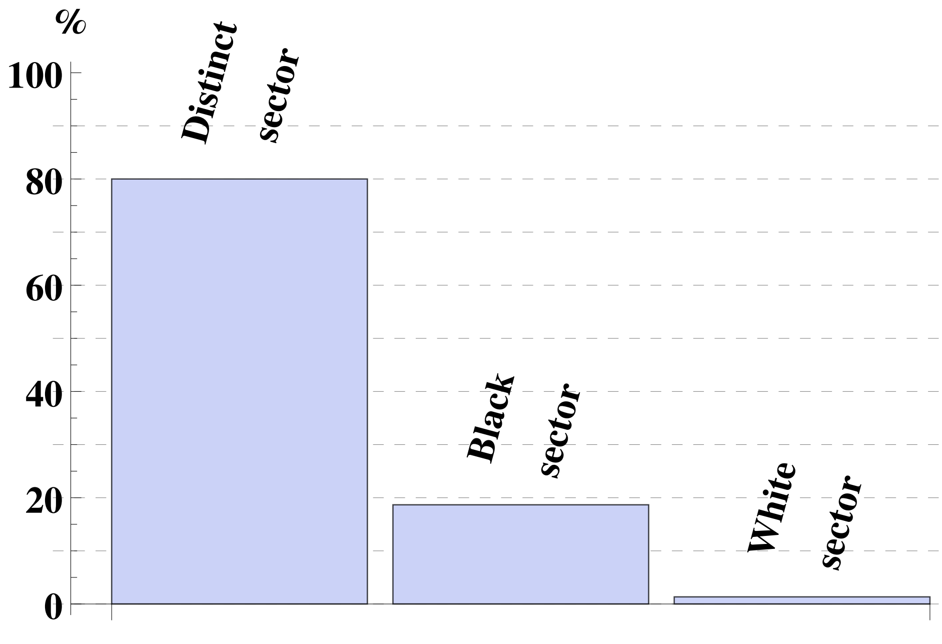

Figure 6: Results of the Coordination Treatment (high-cognition players)

Blume and Gneezy (2010) claim that "beliefs matter" and we test in Hypothesis 2 if there are subjects whose pessimistic beliefs about the other players' skills lead to coordination failure.

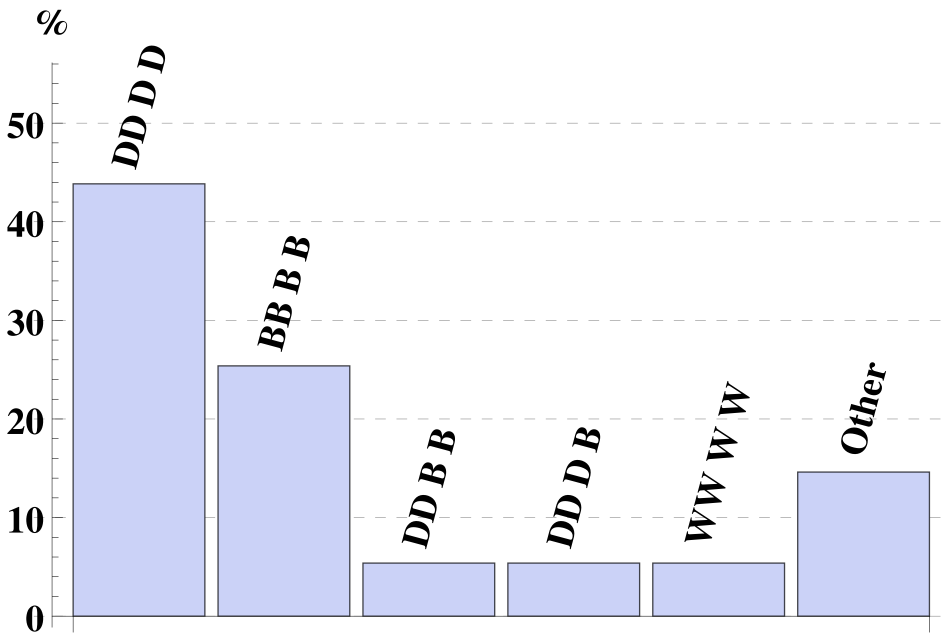

Figure 7: Used strategies

Hypothesis 2 (First-order beliefs matter)**.** *There are high-cognition subjects who choose a black sector in the Prediction Treatment and Cooperation Treatment, i.e., play "DD B B".*

Our data confirms this hypothesis. Figure 7 shows us the results of all players, Figure 8 of the high-cognition players. In these figures we can see that about 9% of the high-cognition players (or 5% of all players) have a first-order belief problem, leading to coordination failure. As the fraction of strategies leading to this conclusion is very small (0.64%) we can reject the null hypothesis that this result is due to chance ($p < 0.00001$).

But do players really use higher-order beliefs in this type of games? Hypothesis 3 tests for this question.

Hypothesis 3 (Higher-order beliefs matter)**.** *There are high-cognition subjects who play the distinct sector in the Prediction Treatment and a black sector in the Cooperation Treatment, i.e., play "DD D B".*

From Figure 7 and Figure 8 we can see that there are high-cognition players who think that their partner is with a high probability of the high type, they, however, still think there are coordination problems. Again, we can reject the null hypothesis at the 1% level ($p < 0.00001$).

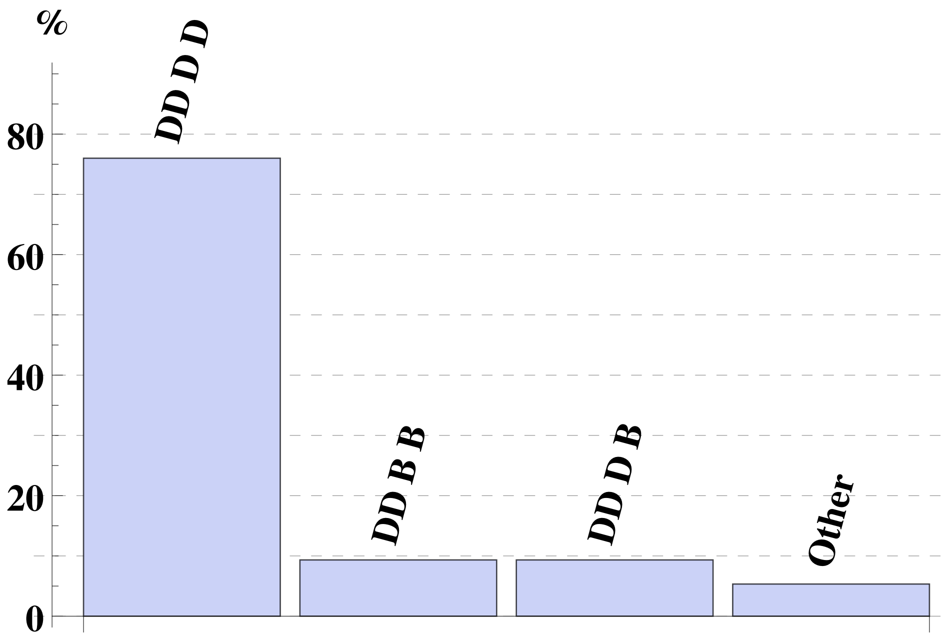

Figure 8: Used strategies (high-cognition players)

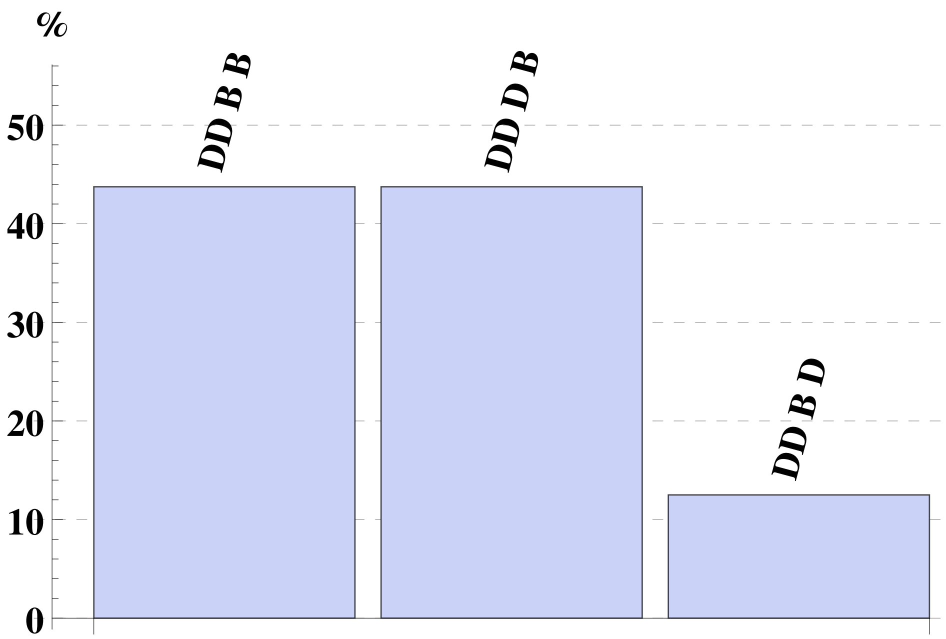

Figure 9: Robustness check

Hypothesis (Robustness check)**.** *"DD B D" is played less often than "DD B B" and "DD D B".*

All these results are statistically significant at the 1% level, however, our design allows for another robustness check: There is a strategy which should not be played by rational players: "DD B D", which is about as likely to be picked at random as ”DD B B“ and ”DD D B“ but can not be explained by our model. Figure 9 shows us that only 2 subjects have chosen this strategy.

We expected to have significant order effects, as in Blume and Gneezy (2010). However, it turns out, that the only robust order effect is a weak effect in the Self Treatment (i.e., more subjects have been able to choose the distinct sector twice later in the experiment).[15] We attribute this to a small change in design. We have explained every treatment before the experiment started and we have conducted a quiz (see the Online Appendix(`http://www.kuelpmann.org/research/`), testing if the instructions have been understood. This probably lead to ”Eureka learning“ before instead of during the experiment.

5. Equilibrium selection and models of higher-order beliefs

In this section we are going to discuss how different models of beliefs and frequently used assumptions on the structure of higher-order beliefs influence the specific game we analyze.

Using the results from the literature on focal points in coordination games (as discussed in Section 7), we know that we can restrict our attention on the two actions with the highest payoffs a_{m-1} and a_m. This simplifies the game to a Bayesian game with two types, a low type whose only attainable action is a_{m-1} and a high type, who has access to a_{m-1} and a_m, without common knowledge about the type distribution. Then, we can denote, with a small abuse of notation, the strategy of a player as the action she chooses if she is of the high-type, i.e., a_m or a_{m-1}, knowing that she will play a_{m-1} if she is of the low type.

Let us first start with the most common assumption, that the *distribution of types is common knowledge*. Then the expected utility of a (risk neutral) high-cognition player is as depicted in Table 2, given her and her partners strategies.[16] p denotes here the probability of a player being of the high type. We can see that the prediction of the model then depends on p. If the probability of a player being of the high type p is too low ($p < \frac{x_{m-1}}{x_{m-1}+x_m}$), only (a_{m-1}, a_{m-1}) will be an equilibrium. In this paper we are going to assume that $p \geq \frac{x_{m-1}}{x_{m-1}+x_m}$ which makes sure that the ”better” equilibrium always exists.[17] For risk-averse players, it is required that

[15] For the full analysis of order effects see Appendix B.

[16] In the analysis we restrict our attention to risk-neutral players. However, the analysis for the case of risk-averse players is analogous and the experimental results are valid for every possible degree of risk aversion.

[17] In the experiment this assumption requires $p > \frac{1}{3}$. As the fraction of high-cognition players is 58%, this assumption is not problematic.

$p \geq \frac{u(x_{m-1})}{u(x_{m-1})+u(x_m)}$, so we know that as long as $p \geq \frac{1}{2}$ the high-type equilibrium always exists, independently of the degree of risk aversion. Furthermore, if the equilibrium exists, it is payoff dominant.

	a_{m-1}	a_m
a_{m-1}	x_{m-1}, x_{m-1}	$(1-p)x_{m-1}, 0$
a_m	$0, (1-p)x_{m-1}$	px_m, px_m

Table 2: Expected utilities of two high-cognition players

Therefore, the prediction of assuming that the *distribution of types is common knowledge* is that, for a high-enough p, we should expect full cooperation.

Monderer and Samet's (1989) *common p-belief* is a generalization of the concept of common knowledge and generates, in this model, the same predictions as assuming that the distribution of types is common knowledge, given a high-enough p.

The game we are analyzing is very close to the original description of a global game as introduced by Carlsson and Van Damme (1993). Written down as in Table 2 it is a very similar game as the main example used in Carlsson and Van Damme (1993). Therefore, we know that, given $\frac{x_{m-1}}{x_m} \leq p \leq \frac{2x_{m-1}}{x_m+x_{m-1}}$ (i.e., (a_m, a_m) is still a Nash equilibrium but (a_{m-1}, a_{m-1}) is risk dominant), (a_{m-1}, a_{m-1}) will be the only rationalizable solution to the global game. Furthermore, Hellwig (2002) shows that higher-order uncertainty about preferences leads to results similar to Carlsson and Van Damme's (1993) higher-order uncertainty about payoffs, i.e., coordination on the "less risky" equilibrium.

Rubinstein (1989) shows that truncating common knowledge at any finite level is equivalent to the situation without any common knowledge at all and therefore suggests that players choose the save strategy a_{m-1}.

Weinstein and Yildiz (2007a) establish a condition, called "global stability under uncertainty" which implies that the change in equilibrium actions is small in the change of kth-order beliefs and higher. Therefore, under this condition, equilibria can be approximated by the equilibrium with lower-order beliefs. Unfortunately, pure coordination games do, in general, not fulfill the conditions for "global stability under uncertainty" as the best responses are very sensitive to every order of beliefs and even a small change in some higher-order belief might make a player change from a_m to a_{m-1}.

Model	Coordination	First-order belief coordination problems	Higher-order belief coordination problems
Common knowledge	Full coordination	No	No
Common p-belief	Full coordination	No	No
Global games	No coordination	Yes	Yes
Almost common knowledge	No coordination	No	Yes.

Table 3: Models of higher-order beliefs

6. The role of beliefs in the fight against female genital mutilation

Female genital mutilation (FGM) is a wide spread problem in many parts of the world and it is estimated to effect up to 200 million women in 2016 (UNICEF (2016)).[18]

The root of the problem seems to be two-fold: A strong strong beliefs in the benefits of it (e.g., hygiene issues, pleasure of the man, ...) and, even among families who do not think it is beneficial, peer pressure.

In the last couple of years, different NGOs and governments have fought against FGM by educating families about the dangers of it. However, it was commonly observed that, while informing families changes their opinion about FGM, the effect in behavior (i.e., not infibulating their daughters) is much smaller than the change in believes (see for example, Bicchieri (2005)).

One possible reason for that might be higher-order beliefs: If I know FGM is bad, but enough other women believe in it, my daughter will not get married if she is not infibulated. Or, even worse, even if everyone knows FGM is bad, but everyone things, that everyone else thinks that infibulation is good, I will still infibulate my daughter.

Here, like in our experiment, coordination on the better "No-FGM" equilibrium could therefore be hindered by pessimistic first- or higher-order beliefs. Even if everyone thinks FGM is bad but thinks that every other family expects it, every family would choose to infibulate her daughters.

The link from this (simplified) problem to our experiment is clear: In both, we have a a bad ("FGM", BB) and a good equilibrium ("no-FGM" and DD) and the latter can only be played by "informed" or "educated" players (i.e., families who

[18]For more information on FGM see for example Gupta (2013) and Bicchieri (2005)

know about the dangers of FGM/players who see the distinct sector). Unfortunately, the lack of common knowledge makes the problems much more complicated and there is a chance that even informed players choose the bad equilibrium due to 1st or higher-order believes.

We have learned from our experiment that, even in simple coordination game in which most subjects were of the "educated" type (i.e., the ones who were able to find the distinct sector), some educated subjects switched to the bad equilibrium because of first- and higher-order believes.

Therefore, it is reasonable to assume that this result might generalize to more general populations and that believes might play a role in the fight against FGM.[19]

What can we now learn from these results for the fight against FGM?

First of all, making the education common knowledge, e.g., by gathering women from the area and explaining everything to everyone at the same time. [20]

Unfortunately, this is not always possible and it might be prohibitively costly. However, due to the structure of the problem (i.e., that the decision for the daughters and sons of the family are separable), we can replicate the Guessing Treament in our experiment. This would reduce the problem from a higher-order belief problem to a first-order belief problem, which is much easier to handle.

To do this, one has to convince every family to (A) not only not infibulate their daughters but also to marry their sons to not infibulated women. And (B) make every family sign a letter of intent that they will marry their sons also to uninfibulated women and will not infibulate their daughters and distribute these letters afterwards to every family.

Then, due to (A), the structure is very similar to our Guessing Treament: now, the decision only depends on the type (i.e., has the other family accepted that FGM is bad) and not on the action of other families (i.e., will they let their son marry my daughter, even if she isn't infibulated). Due to (B) we have made sure that the everyone knows the first-order believe of everyone else.

This inexpensive and simple addition might help to boost success rates of the widely used education approach in the fight against FGM more efficiently.[21]

[19]It might be far-fetched to extrapolate the behavior of students at a British university to villagers in rural Africa or South-East Asia. However, we believe that our results are a good enough reason to either test these in a field experiment or, as an experiment itself, implement the suggestions below into existing programs, as they are very cheap.

[20]This is done by Tostan, an NGO, which emphasizes that education together with public discussion and public declaration is the best way to fight FGM (Tostan (2016)). Their claims are supported by their success rates, which are reported for example in World Bank Group (2012)). While there are probably also other factors at work (,e.g., group dynamics, ...), this also points towards the possible importance of higher-order beliefs.

[21]Of course, as mentioned above, some testing in the field is needed.

7. Related works

There is a large theoretical literature, beginning with the seminal paper on the "email game" by Rubinstein (1989), showing that higher-order beliefs play a role in determining the outcome of a game. For instance, Carlsson and Van Damme (1993) use higher-order beliefs (in their model of global games) to identify the risk-dominant equilibrium as the unique rationalizable outcome of the coordination game. This uniqueness result spawned a large applied literature on, among other areas, bank runs and arms races, in e.g. Morris and Shin (1998), Morris and Shin (2004), Baliga and Sjöström (2004), Corsetti, Dasgupta, Morris, and Shin (2004), and Goldstein and Pauzner (2005). Weinstein and Yildiz (2007b), however, have shown that this uniqueness result, that this whole literature depends on, is fragile to the exact specification of the higher-order belief model. Other "nearby" higher-order belief models have very different "unique" predictions. In fact, they show that any rationalizable outcome of the original game, can be obtained as the unique rationalizable strategy profile of some higher-order belief model.

Weinstein and Yildiz (2007a) establish a condition, called "global stability under uncertainty". This condition implies that, if the change in equilibrium actions is small in the change of kth-order beliefs and higher, equilibria can be approximated by the equilibrium with at most kth-order beliefs. Unfortunately, pure coordination games do not fulfill "global stability under uncertainty".

Strzalecki (2014) and Kneeland (2016) develop different non-equilibrium approaches, inspired by the experimental literature discussed later, using bounded levels of reasoning to explain behavior in coordinated attack problems (e.g. Rubinstein's (1989) email game).
A more in-depth discussion of models of higher-order beliefs and their predictions of the results of our experiment can be found in Section 5.

The experimental literature, however, has so far mostly focused on strategic uncertainty. The most prominent example for this is probably the literature on level-k thinking or cognitive hierarchy models, which was started by Nagel (1995) and Stahl and Wilson (1995). In recent years, there have been many studies conducted, using and analyzing level-k reasoning, for example Ho, Camerer, and Weigelt (1998), Costa-Gomes, Crawford, and Broseta (2001), Camerer, Ho, and Chong (2004) and Crawford, Gneezy, and Rottenstreich (2008). For a recent survey, see Crawford, Costa-Gomes, and Iriberri (2013).

But there also have been works which do not focus on strategic uncertainty. For example Heinemann, Nagel, and Ockenfels (2004), Cornand (2006), Cabrales, Nagel, and Armenter (2007) and Duffy and Ochs (2012) who directly test implications of the theory of global games, i.e. individuals play an incomplete information game as in Carlsson and Van Damme (1993). The results however, are mixed and range from full support to full rejection of the predictions made by global games.

Another, closely related work is Kneeland (2015), in which she explores the level of rationality, a requirement for higher-order beliefs, of players experimentally. She shows that, in her experiment, 94% of all players are rational with decreasing numbers for second- (71%), third- (44%) and forth-order (22%) rationality.

We explore experimentally the "depth of reasoning" individuals employ when playing slightly difficult coordination games. In fact we want to abstract away from purely strategic concerns by only looking at coordination games in which the incentives of the players are perfectly aligned and a Pareto-dominant equilibrium exists. The fundamental uncertainty in the model will be one about the cognitive abilities of the opponents.

Differences in cognitive abilities have been studied before, for example by Gill and Prowse (2016), who have shown that more cognitively able subjects converge, in repeated p-beauty contests, more frequently to equilibrium play and earn more. Furthermore, Proto, Rustichini, and Sofianos (2014) have shown that intelligence affects the results of repeatedly played prisoner's dilemmas, in which groups of higher intelligence tend to cooperate more frequently in later stages of the game. Agranov, Potamites, Schotter, and Tergiman (2012) have shown, by manipulating the perception of the cognitive levels of other players, that beliefs about the level of reasoning do play a significant role in the presence of strategic uncertainty. Alaoui and Penta (2015) establish a framework in which the depth of reasoning is endogenously determined by different cognitive costs of reasoning.

The way we model cognitive differences however, builds on another branch of literature. Motivated by Schelling's (1960) discussion of focal points, a variety of authors have tried to formally capture his ideas, most notably Bacharach (1993) and Sugden (1995). The importance of focal points is supported by many experiments, for example by Mehta, Starmer, and Sugden (1994), who have replicated Schelling's results and have shown that coordinating on a focal point is different from accidental coordination. Crawford, Gneezy, and Rottenstreich (2008) have shown that, in a pure coordination game with symmetric payoffs, salient labels lead to a high percentage of coordination whereas even slight asymmetries in payoffs might lead to a coordination failure. Isoni, Poulsen, Sugden, and Tsutsui (2013) extend the analysis to bargaining problems and show that payoff-irrelevant clues help to improve coordination, even if there is no efficient or equal division.

In the absence of clues however, the theory of focal points can not be applied. Formally the absence of clues can be modeled as symmetries between strategies and players in a given game. In fact Nash (1951) has already discussed equilibrium under symmetry restrictions (and shown existence also of such symmetric (mixed) equilibria for finite games). Crawford and Haller (1990) have defined symmetries in games and used these definitions to see what focal points in highly symmetric repeated coordination games would look like.[22] Blume (2000) has further developed this symmetry concept to talk about play under the absence of a common language. Other notions of symmetries have been put forward and studied in Harsanyi and Selten (1988), Casajus (2000) and Casajus (2001). Alós-Ferrer and Kuzmics (2013) have then clarified the difference between different notions of symmetries and characterized all the possible ways a frame (the way a game is presented to players in the lab, for instance) could lead to different symmetry restrictions (and therefore to different focal points).

All these models of symmetries and restrictions are implicitly or explicitly investigated under the assumption of perfectly rational individuals. However, identifying all symmetries (and especially non-symmetries) in a game can be a difficult task. Bacharach (1993) has proposed his variable frame theory to allow for individual players with different states of mind or, as developed by Blume (2000) and employed by Blume and Gneezy (2000) and Blume and Gneezy (2010), with different cognitive abilities.

This finally brings us to the goal of our study. We want to take up the experimental results and setup of Blume and Gneezy (2010) to analyze the effects of higher-order beliefs. They were able to show that participants form beliefs about the cognitive abilities of other participants and, if these beliefs are pessimistic, they hinder coordination between the players. However, they have not taken into account the effect of higher-order beliefs about cognitive abilities. Therefore, we modify their experimental setup in order to distinguish the effect of first-order beliefs players form about the cognitive ability of their opponents and higher-order beliefs.

8. Conclusion

We have seen that, in this game, absence of common knowledge was not enough to prevent subjects to choose the Pareto-dominant equilibrium strategy, as 76% of the high-cognition players have done so. However, we still have a fraction of players who have beliefs that lead to coordination failure (around 18%) and of these only half could be attributed to first-order beliefs.

Of the models of higher-order beliefs discussed in Section 2.2 and Section 5, only

[22] Bhaskar (2000) and more comprehensively Kuzmics, Palfrey, and Rogers (2014), have studied theoretically and in the latter case also experimentally, what the possible focal points of the symmetric repeated battle-of-the-sexes and its generalizations could be.

"assuming common knowledge"[23] or a common p-belief among high-cognition players about the type distribution were able to explain coordination on the payoff-dominant equilibrium. However, these assumptions can not explain any coordination failure due to beliefs, as the beliefs are fixed by the model, whereas the models which can explain this type of coordination failure predict playing the payoff-dominated strategy.

Therefore, as we have observed a coordination rate of about 76%, assuming common knowledge (or a common p-belief among high-cognition players) might be the best tractable approximation available in coordination games without common knowledge, depending on the focus of the research.

Our work opens up some questions for future research: Can these results be generalized to other populations and environments? Are there certain parts of the populations who are more likely to exhibit first- or higher-order beliefs which lead to coordination failure? Are there other, maybe easier methods to make something common knowledge? Furthermore, it might be worthwhile to check more general structures of higher-order beliefs or if non-equilibrium models like Strzalecki (2014) or Kneeland (2016) can explain this phenomenon better.

[23] Meaning that one assumes that high-cognition players have common knowledge about the type distribution among themselves.

Appendix A Belief hierarchies

Let $B_i^0 := T_j$ and $B_i^k = T_j \times \Delta(B_i^{k-1})$ with $\Delta(B)$ being the space of probability measures on B and $\Delta(X)$ being the space of probability measures on the Borel field of X, endowed with the weak topology. Using this notation, we can define a belief hierarchy as follows.

Definition 1 (Belief hierarchy)**.** A *k-th order belief* is defined as

$$b_i^k \in \Delta(B_i^k)$$

with $B_i^0 = T_j$ and $B_j^k = T_j \times \Delta(B_j^{k-1})$
Furthermore, let us set $b_i^0 := t_i$.
A belief hierarchy of a player i is then $b = \{b_i^0, b_i^1, \dots .\}$

We therefore have a first order belief $b_i^1 \in \Delta(\{low, high\}) = [0,1]$ and higher-order beliefs $b_i^k \in [0,1]^k$.

Furthermore, we assume these beliefs to be coherent, i.e. that beliefs of different orders do not contradict one another,[24] and that a low-cognition type does not know about higher cognitive types, i.e., $b_i^k = 0 \Rightarrow b_i^{k+1} = 0 \quad \forall k \geq 0$.

This excludes, on the one hand, that a low-cognition player thinks that the other player is a high-cognition player and, on the other hand, that a player has a first-order belief that the other player is of a the high type and a higher-order belief that the player is of the low type.

Appendix B Order effects

We have briefly discussed order effects in Section 4. The two main effect of order effects are that (a) the order of treatments had no effect on our main results and (b) the only significant order effect is that, that the number of mistakes in the Self is getting smaller, if the treatment is conducted later.

The most important order effect, we expected in this experiment was "Eureka!" learning, i.e., "Having a player play against himself may trigger an insight that switches a player from low to high cognition ("Eureka!" learning). There may be an uncertainly principle at work here in that we cannot measure a player's cognition without altering it." (Blume and Gneezy (2010)).

Thus, we implemented a random ordering of treatments on the subject level, to be able to control for this.[25]

[24] I.e., higher-order beliefs of a player mapped onto the space of beliefs of a lower order are the same.
[25] We randomized on the subject level to prevent session effects. Furthermore, we thought that randomizing the order for all 130 subjects individually would guarantee a somewhat balanced distribution

Treatment	Self			Prediction			Coordination		
Order	**DD**	**BB**	**Other**	**D**	**B**	**W**	**D**	**B**	**W**
2nd and 3rd	0.40	0.13	0.08	1.00	0.65	0.44	0.25	0.16	0.39
1st and 3rd	0.67	0.18	0.0004	0.29	0.66	0.26	0.02	0.10	0.39
1st and 2nd	0.17	1.00	0.10	0.27	0.26	1.00	0.16	0.84	0.02

Table 4: Order effects of the different treatments

First of all, let us have a look at (a), i.e., the possible effects of the order of treatments on our main results. For certain orders of treatments, "Eureka!" learning might have been misidentified for either first- or higher-order beliefs.

If "Eureka" learning occures between the Prediction Treatment and the Coordination Treatment, we would misidentify this person as someone with higher-order believe problems (i.e., she would play DD D B). However, this is not a problem, as, of our 7 higher-order believe subjects, 4 played the Coordination Game before the Prediction Game (compared to 50.1% of all subjects).

Misidentifying can a subject as having a first-order belief problem (DD B B) due to an order effect can only happen if the Self Treatment is played last. Of the 7 subjects who we identify as having a first-order belief problem only 3 subjects played the Self Treatment last (compared to 40.1% of all subjects - this is higher than the expected 33% due to the unbalanced randomization).[26]

Therefore, we believe that order effects had no significant effect on the main results.

However, having a look at the general order effects might be interesting nonetheless. Therefore, we had a look at every possible order effect and tested every possible of the 27 combinations (3 treatments, 3 choices each and 3 combinations of orderings for the three treatments).

We report the (rounded) p-values of a two-sided Fisher's test in Table 4. In it, we test if there are significant differences in the frequency of a certain action (e.g. picking the distinct sector D), when it is played first against when it is played second (and every other combination).[27]

Let's have a closer look at every treatment:

In the self treatment, one row stands out: The people who have chosen something "Other", i.e., the people who made a mistake. Here, we can see that there is

of orders. It turns out z-Tree does not agree, thus we now have a somewhat unbalanced order distribution ranging from 15 to 32 subjects in each of the 6 groups. Fortunately, this does not effect the results of the experiment.

[26] To see the complete data, the calculations and some additional robustness checks, please refere to the Online Appendix and the R code in "OrderEffects.R" (`http://www.kuelpmann.org/research/`).

[27] In the Self Treatment, we have only disinguised between DD, BB and everything else, as everything else was rare enough (11 subjects overall). In the following we consider "Other" as subjects making a mistake of some kind.

a significant order effect, especially between the Self Treatment being played first and last. The direction is as expected: People are getting better at playing against themselves and, the effect is the strongest if we compare the self treatment being played first and last.

Apart from this, it seems that there are no order effects in the Self Treatment.[28]

In the Prediction treatment, we don't see any significant order effects in either direction.

In the Coordination treatment, we see three two effects: Comparing the 1st and 3rd round, we see significantly more subjects playing *D* and less *B* in the 3rd round than in the 1st round. However, the difference between the 1st and 2nd and the 2nd and 3rd is not significant.

Furthermore, a see a significant lower amount of subjects playing W if the Treatment is played first than second.

Thus, while we see a strong "Eureka" learning effect in the Self Treatment, the "Eureka" learning effect vanishes in almost every other treatment.[29]

Why did Blume and Gneezy (2010) encounter strong "Eureka!"-learning effects whereas we had (almost) no significant effect. We attribute this to the fact that the participants were instructed in all three treatments before they played the first game which most likely triggered the learning before the first decision, whereas in Blume and Gneezy (2010) the instructions for the second treatment were distributed after completion of the first treatment. Furthermore, we used more extensive instructions and a quiz to make sure the instructions where understood.

[28]That we don't see any significant change due to subjects moving from the "Other" group into one of the other two groups is not surprising, as only 11 out of 130 subjects have been clasified as "Other"in the Self Treatment.

[29]We do not want to claim the existence or non-existence of any order effects here, as we clearly not have enough data to show non-existence and a p-value of 0.02 in two of 27 tests is no sufficient basis for the existence. The only thing we claim is that there is an order effect in the Self Treatment (less mistakes later on) and, if there is an order effect anywhere else, it is not particulary strong.

Appendix C Data

For the complete data set please refer to the Online Appendix (`http://www.kuelpmann.org/research/`).

Self		Guessing	Coordination	Self		Guessing	Coordination	Self		Guessing	Coordination
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	D	B	B	B	B
D	D	D	D	D	D	D	B	B	B	B	B
D	D	D	D	D	D	D	B	B	B	B	B
D	D	D	D	D	D	D	B	B	B	B	B
D	D	D	D	D	D	D	B	B	B	B	B
D	D	D	D	D	D	D	B	B	B	B	B
D	D	D	D	D	D	D	B	B	B	B	B
D	D	D	D	D	D	D	B	B	B	B	B
D	D	D	D	D	D	B	D	B	B	B	B
D	D	D	D	D	D	B	D	B	B	B	B
D	D	D	D	D	D	B	B	B	B	B	B
D	D	D	D	D	D	B	B	B	B	B	B
D	D	D	D	D	D	B	B	B	B	B	B
D	D	D	D	D	D	B	B	B	B	B	B
D	D	D	D	D	D	B	B	B	B	B	B
D	D	D	D	D	D	B	B	B	B	B	B
D	D	D	D	D	D	B	B	B	B	B	W
D	D	D	D	D	D	W	D	B	B	B	W
D	D	D	D	D	D	W	W	B	B	W	D
D	D	D	D	D	B	D	D	B	B	W	B
D	D	D	D	B	B	D	D	B	W	D	D
D	D	D	D	B	B	D	D	B	W	B	B
D	D	D	D	B	B	D	B	W	B	W	D
D	D	D	D	B	B	D	B	W	W	D	D
D	D	D	D	B	B	D	B	W	W	D	D
D	D	D	D	B	B	D	B	W	W	B	D
D	D	D	D	B	B	D	B	W	W	W	B
D	D	D	D	B	B	B	B	W	W	W	W
D	D	D	D	B	B	B	B	W	W	W	W
D	D	D	D	B	B	B	B	W	W	W	W
D	D	D	D	B	B	B	B				
D	D	D	D	B	B	B	B				

Table 5: Strategic choice data

References

Marina Agranov, Elizabeth Potamites, Andrew Schotter, and Chloe Tergiman. Beliefs and endogenous cognitive levels: An experimental study. *Games and Economic Behavior*, 75(2): 449–463, 2012.

Larbi Alaoui and Antonio Penta. Endogenous depth of reasoning. *The Review of Economic Studies*, 2015. doi: 10.1093/restud/rdv052.

Carlos Alós-Ferrer and Christoph Kuzmics. Hidden symmetries and focal points. *Journal of Economic Theory*, 148(1):226–258, 2013.

Ayala Arad and Ariel Rubinstein. The 11–20 money request game: a level-k reasoning study. *The American Economic Review*, 102(7):3561–3573, 2012.

Michael Bacharach. *Variable Universe Games*. MIT Press, 1993.

Sandeep Baliga and Tomas Sjöström. Arms races and negotiations. *The Review of Economic Studies*, 71(2):351–369, 2004.

V. Bhaskar. Egalitarianism and efficiency in repeated symmetric games. *Games and Economic Behavior*, 32(2):247–262, 2000.

Cristina Bicchieri. *The grammar of society: The nature and dynamics of social norms*. Cambridge University Press, 2005.

Andreas Blume. Coordination and learning with a partial language. *Journal of Economic Theory*, 95(1):1–36, 2000.

Andreas Blume and Uri Gneezy. An experimental investigation of optimal learning in coordination games. *Journal of Economic Theory*, 90(1):161–172, 2000.

Andreas Blume and Uri Gneezy. Cognitive forward induction and coordination without common knowledge: An experimental study. *Games and Economic Behavior*, 68(2):488–511, 2010.

Adam Brandenburger and Eddie Dekel. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59(1):189–198, 1993.

Antonio Cabrales, Rosemarie Nagel, and Roc Armenter. Equilibrium selection through incomplete information in coordination games: An experimental study. *Experimental Economics*, 10 (3):221–234, 2007.

Colin F Camerer, Teck-Hua Ho, and Juin-Kuan Chong. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, pages 861–898, 2004.

Hans Carlsson and Eric Van Damme. Global games and equilibrium selection. *Econometrica*, pages 989–1018, 1993.

André Casajus. Focal points in framed strategic forms. *Games and Economic Behavior*, 32(2): 263–291, 2000.

André Casajus. Static focal points in framed strategic games. In *Focal Points in Framed Games*, pages 9–54. Springer, 2001.

Michael Suk-Young Chwe. *Rational ritual: Culture, coordination, and common knowledge*. Princeton University Press, 2013.

Russell W Cooper, Douglas V DeJong, Robert Forsythe, and Thomas W Ross. Selection criteria in coordination games: Some experimental results. *The American Economic Review*, 80(1): 218–233, 1990.

Camille Cornand. Speculative attacks and informational structure: An experimental study. *Review of International Economics*, 14(5):797–817, 2006.

Giancarlo Corsetti, Amil Dasgupta, Stephen Morris, and Hyun Song Shin. Does one soros make a difference? A theory of currency crises with large and small traders. *Review of Economic Studies*, pages 87–113, 2004.

Miguel Costa-Gomes, Vincent P Crawford, and Bruno Broseta. Cognition and behavior in normal-form games: An experimental study. *Econometrica*, pages 1193–1235, 2001.

Vincent P Crawford and Hans Haller. Learning how to cooperate: Optimal play in repeated coordination games. *Econometrica*, pages 571–595, 1990.

Vincent P Crawford, Uri Gneezy, and Yuval Rottenstreich. The power of focal points is limited: Even minute payoff asymmetry may yield large coordination failures. *The American Economic Review*, pages 1443–1458, 2008.

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

John Duffy and Jack Ochs. Equilibrium selection in static and dynamic entry games. *Games and Economic Behavior*, 76(1):97–116, 2012.

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

Sidney D Gamble. The disappearance of foot-binding in Tinghsien. *American Journal of Sociology*, pages 181–183, 1943.

David Gill and Victoria L Prowse. Cognitive ability, character skills, and learning to play equilibrium: A level-k analysis. *Journal of Political Economy*, 2016.

Itay Goldstein and Ady Pauzner. Demand–deposit contracts and the probability of bank runs. *The Journal of Finance*, 60(3):1293–1327, 2005.

Geeta Rao Gupta. Female genital mutilation/cutting: a statistical overview and exploration of the dynamics of change. *Reproductive Health Matters*, 21(42):184–191, 2013.

John C Harsanyi. Games with incomplete information played by bayesian players part i-iii. *Management Science*, 14(5):320–334, 1967/68.

John C Harsanyi and Reinhard Selten. *A general theory of equilibrium selection in games*, volume 1. The MIT Press, 1988.

Frank Heinemann, Rosemarie Nagel, and Peter Ockenfels. The theory of global games on test: Experimental analysis of coordination games with public and private information. *Econometrica*, 72(5):1583–1599, 2004.

Christian Hellwig. Imperfect common knowledge of preferences in global coordination games. *UCLA Economics Online Papers*, 211, 2002.

Teck-Hua Ho, Colin Camerer, and Keith Weigelt. Iterated dominance and iterated best response in experimental" p-beauty contests". *American Economic Review*, pages 947–969, 1998.

Andrea Isoni, Anders Poulsen, Robert Sugden, and Kei Tsutsui. Focal points in tacit bargaining problems: Experimental evidence. *European Economic Review*, 59:167–188, 2013.

Margaret E Keck and Kathryn Sikkink. *Activists beyond borders: Advocacy networks in international politics*, volume 6. Cambridge Univ Press, 1998.

Terri Kneeland. Identifying higher-order rationality. *Econometrica*, 83(5):2065–2079, 2015.

Terri Kneeland. Coordination under limited depth of reasoning. *Games and Economic Behavior*, 2016.

Christoph Kuzmics, Thomas Palfrey, and Brian W Rogers. Symmetric play in repeated allocation games. *Journal of Economic Theory*, 154:25–67, 2014.

Robert Duncan Luce and Howard Raiffa. *Games and decisions: Introduction and critical surveys*. Wiley New York, NY., 1957.

Gerry Mackie. Ending footbinding and infibulation: A convention account. *American Sociological Review*, pages 999–1017, 1996.

Gerry Mackie and John LeJeune. Social dynamics of abandonment of harmful practices: a new look at the theory. *Special Series on Social Norms and Harmful Practices, Innocenti Working Paper*, 6, 2009.

Judith Mehta, Chris Starmer, and Robert Sugden. The nature of salience: An experimental investigation of pure coordination games. *The American Economic Review*, 84(3):658–673, 1994.

Jean-Francois Mertens and Shmuel Zamir. Formulation of bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14(1):1–29, 1985.

Dov Monderer and Dov Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1(2):170–190, 1989.

Stephen Morris and Hyun Song Shin. Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, pages 587–597, 1998.

Stephen Morris and Hyun Song Shin. Coordination risk and the price of debt. *European Economic Review*, 48(1):133–153, 2004.

Rosemarie Nagel. Unraveling in guessing games: An experimental study. *The American Economic Review*, pages 1313–1326, 1995.

John Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.

Martin T Orne. On the social psychology of the psychological experiment: With particular reference to demand characteristics and their implications. *American Psychologist*, 17(11): 776, 1962.

Eugenio Proto, Aldo Rustichini, and Andis Sofianos. Higher intelligence groups have higher cooperation rates in the repeated prisoner's dilemma. *IZA Discussion Paper*, 2014.

Ariel Rubinstein. The electronic mail game: Strategic behavior under" almost common knowledge". *The American Economic Review*, pages 385–391, 1989.

Thomas C Schelling. *The strategy of conflict*. Harvard University Press, 1960.

Dale O Stahl and Paul W Wilson. On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, 10(1):218–254, 1995.

Tomasz Strzalecki. Depth of reasoning and higher order beliefs. *Journal of Economic Behavior & Organization*, 108:108–122, 2014.

Robert Sugden. A theory of focal points. *The Economic Journal*, pages 533–550, 1995.

Tostan. Female genital cutting. http://www.tostan.org/female-genital-cutting, 2016. Accessed: 2016-05-21.

UNICEF. Female genital mutilation/cutting: a global concern. http://www.unicef.org/media/files/FGMC_2016_brochure_final_UNICEF_SPREAD.pdf, 2016.

John B Van Huyck, Raymond C Battalio, and Richard O Beil. Tacit coordination games, strategic uncertainty, and coordination failure. *The American Economic Review*, pages 234–248, 1990.

Jonathan Weinstein and Muhamet Yildiz. Impact of higher-order uncertainty. *Games and Economic Behavior*, 60(1):200–212, 2007a.

Jonathan Weinstein and Muhamet Yildiz. A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica*, 75(2):365–400, 2007b.

World Bank Group. *World Development Report 2012: Gender Equality and Development*. World Bank Publications, 2012.