

An Experiment on Learning in a Multiple Games Environment

Citation for published version (APA):

Grimm, V., & Mengel, F. (2009). *An Experiment on Learning in a Multiple Games Environment*. Maastricht University School of Business and Economics. METEOR Research Memorandum No. 007
<https://doi.org/10.26481/umamet.2009007>

Document status and date:

Published: 01/01/2009

DOI:

[10.26481/umamet.2009007](https://doi.org/10.26481/umamet.2009007)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Veronika Grimm, Friederike Mengel

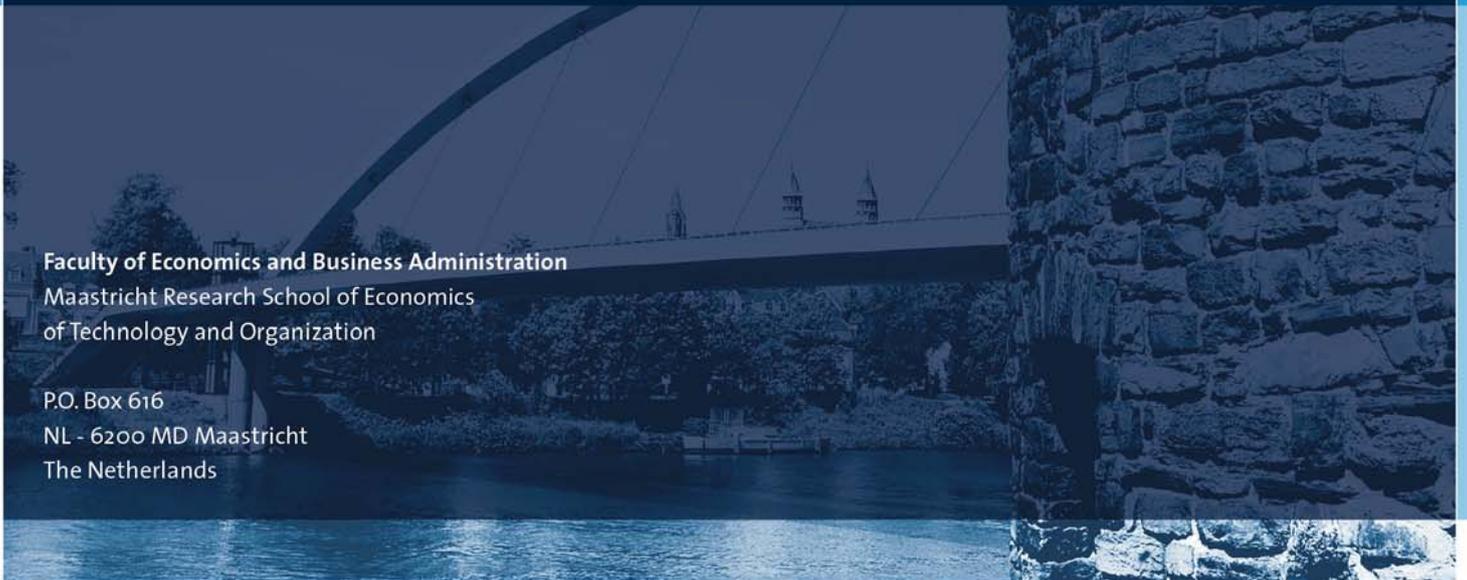
**An Experiment on Learning in a
Multiple Games Environment**

RM/09/007

METEOR

Faculty of Economics and Business Administration
Maastricht Research School of Economics
of Technology and Organization

P.O. Box 616
NL - 6200 MD Maastricht
The Netherlands



An Experiment on Learning in a Multiple Games Environment*

VERONIKA GRIMM[†]

University of Erlangen–Nuremberg

FRIEDERIKE MENGEL[‡]

Maastricht University

February 6, 2009

Abstract

We study experimentally how players learn to make decisions if they face many different (normal-form) games. Games are generated randomly from a uniform distribution in each of 100 rounds. We find that agents do extrapolate between games but learn to play strategically equivalent games in the same way. If either there are few games or if explicit information about the opponent's behavior is provided (or both) convergence to the unique Nash equilibrium generally occurs. Otherwise this is not the case and play converges to a distribution of actions which is Non-Nash. Action choices, though, that cannot be explained by theoretical models of either belief-bundling or action bundling are never observed. Estimating different learning models we find that Nash choices are best explained by finer categorizations than Non-Nash choices. Furthermore participants scoring better in the "Cognitive Reflection Test" choose Nash actions more often than other participants.

JEL-classification: C70, C91

Keywords: Game Theory, Learning, Multiple Games.

*We thank Rene Fahr for a helpful suggestion and seminar participants in Barcelona (Pompeu Fabra) and California-Irvine for helpful comments. We also thank Meike Fischer for very valuable research assistance. Financial support by the *Deutsche Forschungsgemeinschaft* and the *Spanish Ministry of Education and Science* (grant SEJ 2004-02172) is gratefully acknowledged.

[†]University of Erlangen–Nuremberg, Lehrstuhl für Volkswirtschaftslehre, insb. Wirtschaftstheorie, Lange Gasse 20, D-90403 Nürnberg, Germany, Tel. +49 (0)911 5302-224, Fax: +49 (0)911 5302-168, email: veronika.grimm@wiso.uni-erlangen.de

[‡]Maastricht University, Department of Economics (AE1), PO Box 616, 6200MD Maastricht, Netherlands. *e-mail*: F.Mengel@algec.unimaas.nl

1 Introduction

The theory of learning in games has largely been concerned with agents that face one well specified game and try to learn optimal actions in that game. In addition it is often assumed that the description of the game (players, strategies and payoffs) are known by all the players interacting in the game. Reality, though, is much more complex. Agents face many different games and while they typically know their own preferences, they often do not know those of their opponents. In such cases while they will know the "label" of a game, they will not know the full strategic form.

In this paper we investigate experimentally how agents deal with such complex situations. In particular we study which implications these standard assumptions have on learning and convergence to equilibrium in simple normal form games. We are also interested in whether agents tend to categorize games and how such categorizations will look like. In our experiment participants interact in 3×3 normal form games for 100 rounds. The number of games is varied across treatments. Agents know their own payoffs in each game, but not those of their opponents. In addition in some treatments explicit information about the opponent's past behavior is provided, whereas in others it is not. While we vary the complexity of the environment across treatments (by varying the number of games and the amount of information explicitly provided), the games themselves are simple. In particular they are all acyclic and have a unique strict Nash equilibrium that coincides with the maxmin prediction.

We find that participants do extrapolate between games, but learn to play strategically equivalent games in the same way. If either there are few games or if explicit information about the opponent's behavior is provided (or both) convergence to the unique Nash equilibrium generally occurs. Otherwise this is not the case and play converges to a distribution of actions which is Non-Nash. Action choices, though, that cannot be explained by theoretical models of either belief-bundling or action bundling are never observed. Estimating different learning models we find that Nash choices are best explained by finer categorizations than Non-Nash choices. Furthermore participants scoring better in the "Cognitive Reflection Test" (Frederick, 2005) choose Nash actions more often than other participants.

Some other authors have investigated similarity learning or analogy based decisions in experiments. Huck, Jehiel and Rutter (2007) let participants interact in two different normal form games varying the degree of accessibility of information about the opponent's behavior. They show that Jehiel's (2005) equilibrium concept of "analogy based expectations" can rationalize their data. Stahl and Huyck (2002) let participants interact in different 2×2 stag hunt games and demonstrate that similarity plays a role in the participants' decisions. Finally Haruvy and Stahl (2008) let participants interact in 10 different normal form games to see how and whether they extrapolate between similar games. They show that convergence to Nash equilibrium occurs faster if participants have been exposed to a strategically equivalent game before and rationalize the data through Stahl's n-level learning theory. Of course there is also some theoretical literature dealing with the idea of similarity thinking and learning across

games. Most of the theoretical literature on similarity is concerned with one person decision problems. The probably most well known example is case-based decision theory (Gilboa and Schmeidler 1995, 1996).¹ They argue that if agents face a decision problem for the first time they will reason by analogy to similar situations they have faced in the past. They then derive a utility representation for an axiomatization of such a decision rule. The similarity function from case-based decision theory has been applied to games by for example Li Calzi (1995) or Steiner and Stewart (2007). Jehiel (2005) and Mengel (2008) present different concepts of analogy reasoning. In both papers agents partition games into subsets of games they see as analogous. In Jehiel (2005) these partitions are exogenous, whereas in Mengel (2008) they are endogenous. Also the concept of analogy thinking differs in the two papers. In Jehiel (2005) seeing different games as analogous means having the same expectations about the opponent's behavior, implying that different actions can be chosen in games seen as analogous. By contrast in Mengel (2008) seeing two games as analogous always implies choosing the same actions in these games.²

The paper is organized as follows. In Section 2 we describe the experimental design and our research questions. In Section 3 we present the results and section 4 concludes.

2 Experimental Design and Hypotheses

We consider behavior of agents that face many different strategic situations. In particular, agents interact in different games and might not be able to distinguish all of them due to their limited cognitive capacities. Our aim is to gain insights on how behavior in each individual situation is affected by the complexity of the environment (if at all).

An obvious human reaction to complexity is to form categories of situations that are perceived as "similar" or "close" and treat all games in a category in a similar way. It is, however, not at all obvious, (a) which criteria need to be satisfied to establish the perception of closeness and (b) in what way games of the same category are treated similarly. Our experiment is designed to yield insights on both issues:

- (a) We analyze systematically how people categorize games. Possible criteria for choosing categories could be strategic similarity (roughly, games are perceived as close if there is overlap of the supports of their Nash equilibria), or closeness in the payoff space³ (as measured either through matrix norm of the difference

¹A nice paper is also Rubinstein (1988), who offers an explanation of the Allais-paradox based on similarity.

²There is also a related literature relying on automaton theory to induce complexity costs of playing complicated strategies or distinguishing many games. See among many others Abreu and Rubinstein (1988) or Samuelson (2001).

³See e.g. Rubinstein (1988). Steiner and Stewart (2008) and Li Calzi (1995) use such a criterium for games.

of the payoff matrices or the Hausdorff distance of their convex hulls). Of course categories could also be learned (Mengel, 2008).

- (b) Within those categories subjects could have different kinds of "analogy thinking". One possibility is *belief bundling* across games in the same category (e.g. Jehiel, 2005). A second possibility is *action bundling* across the games in the same category (e.g. Mengel, 2008). While under belief bundling agents will have the same beliefs about the opponent's behavior in all the games within a given category, they can still choose different actions in games contained in the same category. This is (by definition) not possible under action bundling. Agents engaging in action bundling need not form beliefs about the opponent at all (e.g. they could be reinforcement learners).

Our experiment is designed in order to distinguish between those possible categorizations as well as the different forms of similar treatment of games within the same category. Of course, we can also find out whether behavior (or part of it) is random in the sense that it is not guided by any of the above mentioned forces.

2.1 The Experimental Design

Let us now provide details on our experimental design. We let 128 participants anonymously interact in different 3×3 normal form games for 100 rounds. In all treatments subjects were split up equally in row and column players. They were randomly re-matched in each round within groups of eight participants. We ran four treatments in total that differed in (a) the number of different games the subjects were confronted with (either two or six) and (b) whether explicit information about the opponent's play was provided.

In all treatments each of the possible games occurred with equal probability in each round. Thus, in the two treatments with few games each of the two possible games occurred with probability $\frac{1}{2}$, while in the treatments with many games, each game occurred with probability $\frac{1}{6}$. The normal form games we chose are given in table 1.

Note that games 4, 5, and 6 are "copies" of games 1, 2, and 3 (payoffs are monotonically transformed by adding either 5 or 10 to all values). In the two treatments with few games (labeled "F"), subjects played games 1 and 2. In the treatments with many games (labeled "M") subjects played all games shown in table 1.

Note that all games have a unique strict Nash equilibrium. In games 2 and 5 this Nash equilibrium is in strictly dominant strategies and games 3 and 6 are solvable by iterated strict dominance. In games 1 and 4 elimination of a weakly dominated strategy is needed to reach the unique Nash equilibrium. All games are acyclic (Young, 1993) and the unique strict Nash equilibrium coincides with the maximin prediction. We designed the games in this way so that a) it is easy to learn the Nash equilibrium if the game is the only game played and b) there are no conflicting predictions of existing theories of learning in a single game.

Game1	<i>a</i>	<i>b</i>	<i>c</i>		Game4	<i>a</i>	<i>b</i>	<i>c</i>
<i>A</i>	20, 20	15, 10	15, 10	,	<i>A</i>	25, 25	20, 15	20, 15
<i>B</i>	10, 15	25, 10	0, 10		<i>B</i>	15, 20	30, 15	5, 15
<i>C</i>	10, 15	15, 35	35, 0		<i>C</i>	15, 20	20, 40	40, 5
Game2	<i>a</i>	<i>b</i>	<i>c</i>		Game5	<i>a</i>	<i>b</i>	<i>c</i>
<i>A</i>	5, 5	15, 20	5, 10	,	<i>A</i>	15, 15	25, 30	15, 20
<i>B</i>	10, 15	5, 25	10, 10		<i>B</i>	20, 25	15, 35	20, 20
<i>C</i>	20, 5	25, 15	15, 10		<i>C</i>	30, 15	35, 25	25, 20
Game3	<i>a</i>	<i>b</i>	<i>c</i>		Game6	<i>a</i>	<i>b</i>	<i>c</i>
<i>A</i>	15, 10	20, 20	15, 15	,	<i>A</i>	20, 15	25, 25	20, 20
<i>B</i>	15, 20	10, 15	5, 10		<i>B</i>	20, 25	15, 20	10, 15
<i>C</i>	20, 5	15, 35	35, 0		<i>C</i>	25, 15	20, 40	40, 5

Table 1: Payoff Matrices.

On the other hand, if multiple games are played and if agents form categories, the theoretical prediction will depend on a) which categories agents use and b) whether they engage in belief bundling or action bundling. More precisely our games are designed in such a way that belief bundling will lead to an increase in B (b)-choices in games 1 and 4 if either the coarsest category or category $\{1,3,4,6\}$ is chosen. Action bundling of games on the other hand predicts an increase in C (b) choices in games 1,3,4 and 6 whenever the coarsest category (including all games) is used and an increase in B (b) choices (by column players) if category $\{1,3,4,6\}$ is used. (We focus on these categories since those are the ones we find empirically). In games 2 and 5 both predict Nash choices irrespective of the category used.

We generated games randomly from a uniform distribution. The following table shows how often each game occurred in the treatments with few games (F) and many games (M).

	1	2	3	4	5	6
F	50	50	-	-	-	-
M	13	13	17	30	15	12

Table 2: Frequencies of the different games.

Throughout the experiment subjects could only see their own payoffs, but not the payoffs of their match. After each round they were informed of their interaction partner's choice. In all treatments subjects received exactly the same information, however, in two of them it was harder to extract the relevant information since it required substantial reasoning and memorizing capacities.

In the treatments with high accessibility of information (labeled "I"), payoff matrices for all the game (containing only the participants own payoffs) were printed in the instructions that were distributed at the beginning of the experiment.⁴ In each round subjects were, moreover, informed about the average action choices of their past interaction partners in the last five rounds. In the treatments with low accessibility of information, the payoff matrices did not appear in the instructions but were only shown to the subjects for a limited time in each round before the software switched to the decision screen (where the payoffs were not displayed any more). There was also no explicit information given on average action choices of the opponent. However, in principle subjects could extract this information from the information provided to them on the action choice of their match in each round. Subjects were not allowed to take notes throughout the experiment in any of the treatments.

In the following the treatments will be labeled as follows:

	few games	many games
explicit information	FI	MI
no explicit information	F	M

Table 3: Treatments.

In all treatments except FI subjects participated in a cognitive reflection test (Fredericks, 2005) immediately after the experiment. We did not provide material incentives for correct answers in the test.

The experiment took place in December 2007 at the Experimental Laboratory in Cologne. All experimental sessions were computerized.⁵ Written instructions were distributed at the beginning of the experiment.⁶ Sessions lasted between 60min (F) and 90min (M and MI) (including reading the instructions, answering a post-experimental questionnaire and receiving payments). Students earned between 12,90 Euro and 21,50 Euro in the experiment (including the show-up fee of 2.50) On average subjects earned a little more than 18 Euros (all included).

2.2 Research Questions

With this experimental design we want to answer the following research questions. First, we are interested in whether agents forms categories of games that they perceive as similar and, if so, how this affects their behavior:

Q1.1 Do participants extrapolate between different games?

⁴Both in the instructions as well as during the experiment, payoff information for both row and column players was presented in exactly the same way.

⁵The experiment was programmed and conducted with the software z-Tree (Fischbacher 2007). Subjects were recruited using the Online Recruitment System by Greiner (2004).

⁶The instructions for treatment MI, translated from German into English, can be found in the Appendix. Instructions for the remaining treatments are available upon request.

Q1.2 To which extent does the degree of extrapolation depend on the complexity of the situation as measured by a) whether or not explicit information about the opponent's behavior is provided and b) the number of games played?

Q1.3 Can we find evidence of either belief bundling or action bundling or both?

Q1.4 Do participants learn to "recognize" strategically equivalent games and behave accordingly?

Second, we will analyze the dynamics of the learning process:

Q2.1 How do participants learn across games?

Q2.2 Do similarity relations evolve over time?

Q2.3 Are the similarity (analogy) relations that participants use transitive?

Q2.4 Are there differences between participants in how they learn across games?

We will look at three different types of evidence: action choices in the different games, estimated learning models and questionnaire data. In Section 3.1 we start by looking at the action choices of the participants in our experiment allowing us to (partly) address our first set of research questions (Q1.1 – Q1.4). In Section 3.2 then we take a dynamic perspective and try to answer the questions (Q2.1 – Q2.4). In Section 3.3 we will try to relate the behavior of the participants to their score in the cognitive reflection test, allowing us to gain more insights into the last question (Q2.4).

3 Results

This section contains our experimental results. We start by describing action choices in section 3.1. Later, in section 3.2 we analyze which kind of categorization and analogy thinking could explain the data best.

3.1 Action Choices

Games 1 and 4

We first consider action choices in Game 1 and Game 4, which are strategically equivalent. Remember that Games 1 and 4 have a unique strict Nash equilibrium which is given by (A,a) and are acyclic (Young, 1993). Table 4 summarizes the share of A, B and C choices during the last 50 rounds of the experiment separately for row (RP) and column (CP) players and for the two games (Game 1 and Game 4).

As we suspected, treatments with "many" games (M and MI) seem to be more difficult for the players than treatments with few games (F and FI), at least in the sense

Game 1 (RP)	A	B	C	(CP)	a	b	c
FI	0.90	0.08	0.02	FI	0.92	0.07	0.00
F	0.96	0.04	0.00	F	0.95	0.03	0.02
MI	0.61	0.30	0.08	MI	0.76	0.24	0.00
M	0.42	0.46	0.09	M	0.52	0.48	0.00
Game 4 (RP)	A	B	C	(CP)	a	b	c
MI	0.83	0.13	0.04	MI	0.88	0.12	0.00
M	0.48	0.46	0.06	M	0.43	0.56	0.01

Table 4: Share of action choices in game 1 and 4 (last 50 rounds).

that the Nash equilibrium action A is chosen less frequently. The share of A -choices is significantly different across all treatments except for F and FI (Mann-Whitney Test, $p < 0.0001$). This is also illustrated in figures 1 and 2 (here, triangles represent treatments with few games, and hollow symbols stand for no explicit information). The tables also illustrate that choices in Game 4 are not very different from Choices in Game 1 in the last 50 rounds. As expected, thus, agents do learn to "recognize" strategic similarity.⁷

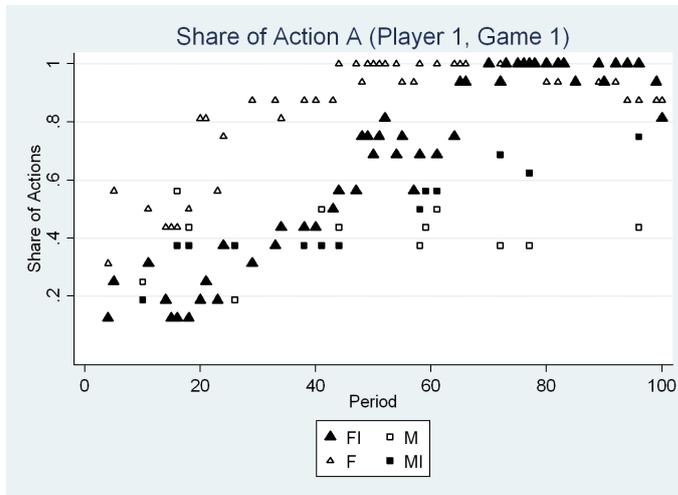


Figure 1: Share of Action A in Game 1 (Player 1) over time

Remember that we designed the games to be acyclic and thus — if learning is not affected by similarity reasoning or other cognitive constraints — participants should converge to Nash equilibrium over time. In fact the best response to the average behavior of the opponents in the whole population (disregarding the fact that each individual was matched only with a subset of all participants) is action A (a) for row

⁷According to distance in payoff space (using either matrix norms or Hausdorff distance) game 2 is for example closer to game 4 than game 1.

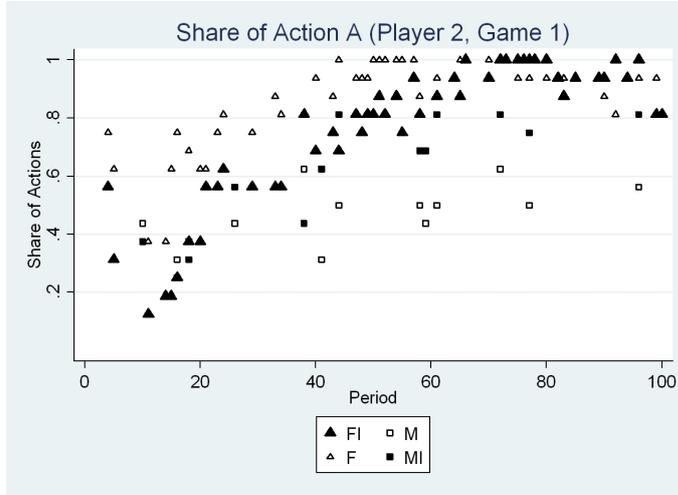


Figure 2: Share of Action A in Game 1 (Player 2) over time

(column) players in treatments FI, F and MI. In treatment M, though, best responses to the actual distribution were given by B (b) for row (column) players respectively, although the expected payoff difference to action A (a) is very small. Figures 1 and 2 illustrate that treatment M is also the only treatment where choices do not converge to the Nash prediction.

How are these facts reflected in the average payoffs of our participants? Table 5 contains information on average payoffs in the last 50 rounds of the experiment. We partitioned the participants according to the action they chose most of the time during the experiment. It can be seen that participants choosing predominantly action C obtain substantially lower payoffs on average than all others. The payoff differences between agents choosing predominantly action A (a) or action B (b) are not significant in any of the treatments. The numbers in brackets are standard deviations.⁸

Game 1	A	B	C	a	b
FI	18.6 (4.2)	17.9 (6.1)	17.5 (6.1)	18.2 (5.0)	15.1 (7.4)
F	18.7 (3.3)	-	-	18.7 (4.1)	-
MI	17.7 (4.0)	18.2 (6.7)	14.2 (3.3)	17.4 (5.2)	17.8 (11.1)
M	16.7 (4.8)	16.7 (6.0)	15.3 (6.4)	16.7 (4.9)	16.2 (9.7)
Game 4	A	B	C	a	b
MI	23.3 (4.3)	24.9 (6.5)	20.6 (3.1)	23.0 (5.6)	22.1 (9.8)
M	22.6 (4.3)	22.5 (6.5)	19.4 (4.3)	21.7 (4.6)	20.8 (9.8)

Table 5: Payoffs by game and action choice (players are classified A (B, C) if they chose that action most of the time)

⁸The last column is missing in the table, as there are no column players who chose mainly action c .

Let us now come back to the question whether action choices eventually converge to Nash equilibrium. Figures 1 and 2 illustrate that in FI and F (black and hollow triangles) action choice indeed converges to the unique Nash equilibrium (A,a) in both, game 1 and game 4. In fact, 90% (92%) of all row (column) players choose action A on average in FI. In treatment F 96% (95%) do so. The only difference between these two treatments is basically a difference in the speed of convergence (See figures 1 and 2). Surprisingly, convergence is slower in treatment FI, where detailed information is being provided. This might be explained by the fact that subjects receive information on their opponents' behavior in the last five rounds, which may slow down convergence.

In treatment M (many games and no explicit information), on the contrary, the distribution of action choices does not seem to converge to equilibrium as figures 3 and 4 illustrate. The figures show how the proportion of A choices in Game 4 evolves for both the row players (denoted "player 1") and the column players ("player 2") in the treatments M (hollow symbols) and MI (black symbols). In treatment MI (many games with explicit information) action choice seems to converge to Nash equilibrium, but convergence is clearly slower as compared to the treatments with few games.⁹ As expected, there is no significant difference between game 1 and game 4. In the more complex environment with many games providing explicit information about the opponent's behavior seems necessary for convergence. Without such information a significant number of choices is Non-Nash, possibly because agents rely on information from several games when making their choices. We will come back to this in Section 3.2.

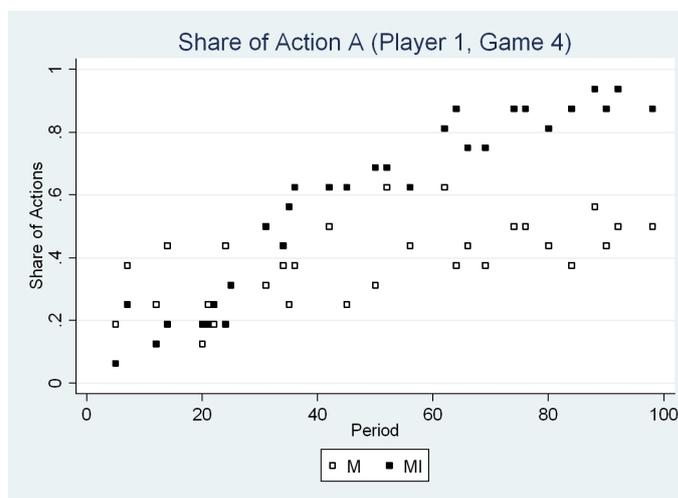


Figure 3: Share of Action A in Game 4 (Player 1) over time

⁹This is still true if one takes into account that agents faced game 1 not as often in the treatments with many games as they did face game 1 in the treatments with few games. Note also that, even though game 4 occurred more than twice as often as game 1 in the treatments with many games, average play is not closer to Nash in game 4 than in game 1.

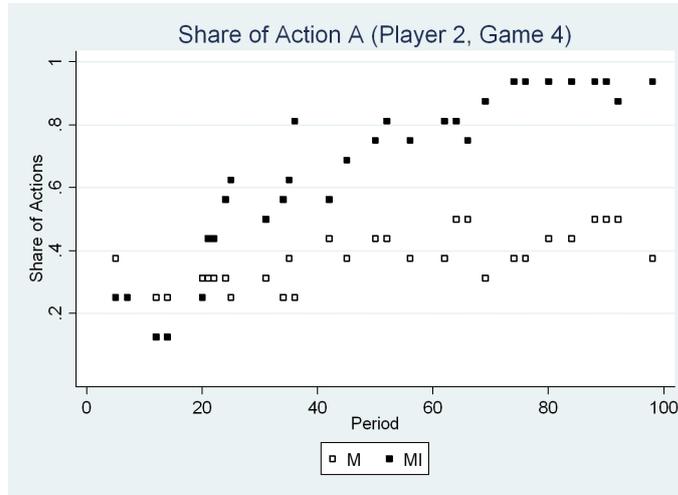


Figure 4: Share of Action A in Game 4 (Player 2) over time

RESULT 1 (ACTION CHOICES GAME 1 AND 4) 1. Action Choices converge to Nash equilibrium (A,a) in treatments FI , F and MI . Convergence is faster in F than in FI and slow in MI . In treatment M there is no convergence to Nash equilibrium.

2. In treatment MI around around 61% (83%) of all choices by row players are A , 30% (13%) B and 8% (4%) C . Choices by column players are 76% (88%) A and 48% (24%) B with no C choices. In treatment M around 42% (48%) of all choices by row players are A , 46% B and 9% (6%) C . Choices by column players are 52% (43%) A and 48% (56%) B with no C choices.

3. Action Choices in Game 4 are not significantly different from Action Choices in Game 1.

Games 2 and 5

In games 2 and 5 both, row and column players, had a strictly dominant strategy (the equilibrium is (C,b)) and - non-surprisingly - they quickly learn to play this strategy. In fact, in none of the treatments and for none of the player roles the overall share of Nash (C,b) choices is below 95%.

RESULT 2 (ACTION CHOICES GAME 2 AND 5) Action choices quickly converge to Nash equilibrium (C,b) in games 2 and 5 in all treatments.

Games 3 and 6

Games 3 and 6 are solvable by iterated elimination of strictly dominated strategies. b.¹⁰ As games 1 and 4 they were acyclic and had a unique strict Nash equilibrium, (A,b). Table 6 shows the distribution of action choices in games 3 and 6 over the last 50 rounds of the experiment.

Game 3 (RP)	<i>A</i>	<i>B</i>	<i>C</i>	(CP)	<i>a</i>	<i>b</i>	<i>c</i>
MI	0.97	0.01	0.01	MI	0.01	0.99	0.00
M	0.54	0.00	0.46	M	0.02	0.98	0.00
Game 6 (RP)	<i>A</i>	<i>B</i>	<i>C</i>	(CP)	<i>a</i>	<i>b</i>	<i>c</i>
MI	0.89	0.01	0.11	MI	0.00	0.99	0.01
M	0.50	0.00	0.50	M	0.01	0.98	0.01

Table 6: Share of action choices in game 3 and 6 (last 50 rounds).

Column players choose strategy b almost always. This is not very surprising since b is a dominant strategy if the column player holds the belief that row players do not choose strategy B. Note that this has to be learned from the row players' behavior since the other player's payoff did not occur in the matrices.

Let us now turn to the results concerning row players. In MI (where it is easy to obtain information about the opponent's play) participants choose the Nash action at least 97% (89%) of the time in game 3(6), but interestingly we observe a lot of C choices (46% and 50%) in treatment M. The share of C-choices is significantly different across the two treatments (Mann-Whitney Test, $p < 0.0001$). This again points to the fact that the environment in treatment M was too complex for subjects to store and process all the information needed for best responses. Again also the distribution of action choices does not differ substantially between games 3 and 6.

Note also that while both belief-bundling and Nash behavior predict A-choices, action bundling predicts an increase in C-choices in these games if agents extrapolate across all games (see also Section 3.2). None of the concepts of analogy reasoning, though, predicts an increase in B (*a,c*) choices, which are choices that are never observed. Behavior thus seems far from random even in the most complex environment (treatment M).

Even though a substantial fraction of row players choose action C, the best response to the average behavior of column players is clearly to choose action A. This is also reflected in the payoffs (see table 7), where those row-players that choose mainly action A earn clearly more than those that choose mainly action C. This suggests that C-choices might be due to the fact that agents choosing mainly C are less willing to engage in costly cognitive reflection and hence use coarser partitions. We will provide some evidence for this in Sections 3.2 and 3.3.

¹⁰B is dominated for the row player, and given the row player does not play B, the column player should play b.

Game 3	<i>A</i>	<i>C</i>
MI	19.0 (2.1)	17.3 (2.5)
M	18.1 (4.2)	15.8 (4.2)
Game 6	<i>A</i>	<i>C</i>
MI	24.1 (2.0)	20.5 (6.4)
M	23.5 (2.3)	20.6 (2.0)

Table 7: Payoffs by game and action choice (players are classified A (B, C) if they chose that action most of the time)

Figures 5 and 6 again illustrate the evolution of action choices A and C over time for row players.

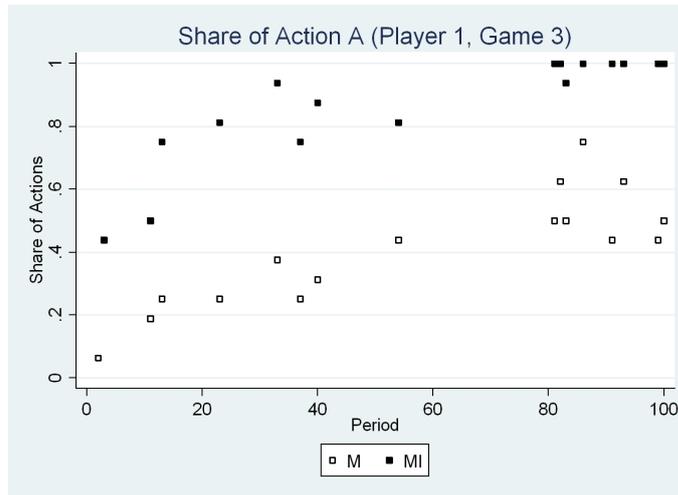


Figure 5: Share of Action A in Game 3 (Player 1) over time

While participants over time learn to choose predominantly action *A* in treatment MI, in treatment M action choices stabilize at a point where roughly half of the choices are *A* and the other half are *C*.

RESULT 3 (ACTION CHOICES GAME 3 AND 6) 1. Column Players choose the action *b* at least 98% of the time in both treatments, M and MI.

2. Row players best respond to this behavior (i.e. play action *A*) at least 89% of the time in MI but only roughly 50% of the time in M.

3. Action Choices in Game 3 are not significantly different from Action Choices in Game 6.

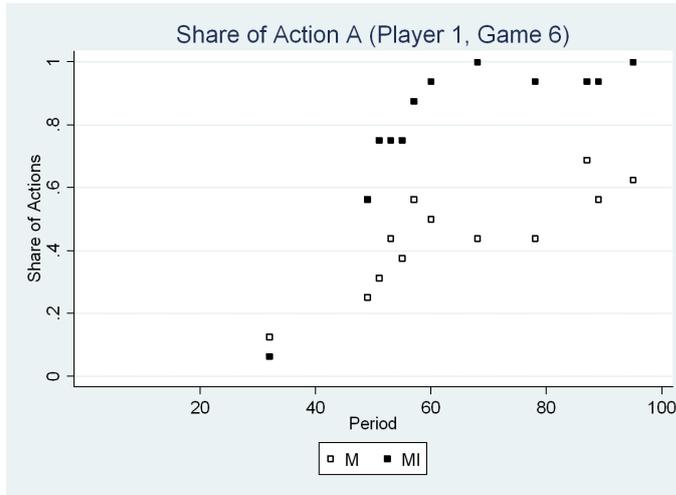


Figure 6: Share of Action A in Game 6 (Player 1) over time

Summary

A first conclusion that we can draw from the descriptive evidence is that (at least for the simple games we chose) participants on average learn to treat strategically equivalent games in the same way. Below we will see that this is also a good description of individual behavior. 70% (80%) of all participants in treatment M (MI) display the same choice behavior in all games that are strategically equivalent.

We have also seen, though, that - unless games have a dominant strategy - participants seem to have different ideas about what the optimal action choice is in each pair of strategically equivalent games. This is most striking in games 3 and 6, where in spite of the fact that the column players choose b almost all the time, some row players think that A is the best choice in these games while others believe that C is the best choice. In addition those participants who choose mainly action C in these games earn clearly lower payoffs due to the fact that C is not a best response to the column players' actions in these games. Also note that the variance is weakly higher for payoffs associated with C choices rather than A choices. Also in games 1 and 4 players seem to have different ideas about optimal choices at least in treatment M. As the share of suboptimal responses declines drastically from M to MI, it seems unlikely that participants have some kind of "psychological" preference for the suboptimal actions. Rather, it seems a reasonable conjecture that these suboptimal choices must have to do with cognitive constraints. In the next section we will study how participants might learn across the different games and whether and how they do categorize games.

Before we do this, though, we will have a quick look at whether it is the same participants that play (or deviate from) the Nash action in the different games. To these ends we again partition the set of all participants according to the action they chose most of the time in a given game. We do this analysis for treatment M, where

we observe most deviations from Nash equilibrium. In this treatment six out of 32 participants choose the Nash equilibrium action most of the time in all games 1 – 6. All others deviate from Nash (most of the time) in at least one of the games. Out of those that deviate in at least one game seven row players choose action C in both games 3 and 6, three (two) row players choose action $B(C)$ most of the time in games 1 and 4 and six column players choose action B most of the time in games 1 and 4. 22 participants choose the same action most of the time in all strategically equivalent games. Two participants play both games 3 and 6 as well as 1 and 4 differently, while eight play only games 1 and 4 differently.

3.2 Categorization and Learning Dynamics

In this section we analyze which categorizations best explain the observed action choices. We moreover investigate whether different "types" of participants (choosing different actions) follow different learning rules. To these ends we estimate two different learning models (a reinforcement learning model and a belief-based model) under different assumptions on the categorizations that players use and then apply the Akaike information criterion (AIC) to evaluate model accuracy.¹¹

More precisely we assume that players are endowed with propensities (or scores) with which they evaluate the different actions. These propensities are updated differently depending on which categorizations the agents hold and on whether they rely on belief-based or reinforcement learning. Finally there will be a choice rule that chooses actions as a function of their propensities.¹² Denote the propensities of player i by $\alpha^i = (\alpha_A^i, \alpha_B^i, \alpha_C^i) \in \mathbb{R}^3$ for the three actions A, B and C and let the payoff of player i in round t be denoted by π^{it} . Denote the category employed by player i by $g^i \subseteq \{1, 2, 3, 4, 5, 6\}$.

Under the reinforcement learning rule (see e.g. Roth and Erev, 1995), if at time t game γ occurs and agent i chooses action $m \in \{A, B, C\}$, propensities are updated as follows:

$$\begin{aligned} \alpha_m^{it+1}(g) &= (1 - \delta)\alpha_m^{it}(g) + \delta\pi^{it} \text{ if } \gamma \in g \\ \alpha_h^i(g) &= (1 - \delta)\alpha_h^{it}(g) \quad \text{for } h \neq m \text{ or if } \gamma \notin g, \end{aligned} \tag{1}$$

where $\delta \in (0, 1)$.¹³ Another model we investigate is belief-based learning as proposed by Sarin and Vahid (1999). According to this rule propensities are updated as follows:

$$\begin{aligned} \alpha_m^{it+1}(g) &= (1 - \delta)\alpha_m^{it}(g) + \delta\pi^{it} \text{ if } \gamma \in g \\ \alpha_h^i(g) &= \alpha_h^{it}(g) \quad \text{for } h \neq m \text{ or if } \gamma \notin g. \end{aligned} \tag{2}$$

¹¹The Akaike criterion is given by $LL-k$ where LL is the log-likelihood and k the number of free parameters of the model. For small sample sizes sometimes the corrected Akaike criterium is used (See Burnham and Anderson, 2002). Since all our competing models have the same number of free parameters we can essentially evaluate them through the log-likelihood, though.

¹²See e.g. Camerer and Ho (1999), Rustichini (1999), or Hopkins (2002, 2007).

¹³See also Mengel (2008).

The difference between these two updating rules is in how actions (or analogy classes) not used are treated. The second rule is typically interpreted as a belief-based rule, as the value of $\alpha_m(g)$ is an estimate of the average payoff of action m in the games contained in category g . In contrast, under the first rule the propensities of action/category pairs, that have not been used for some time decrease with a factor $(1-\delta)$. The $\alpha_m(g)$ in this case cannot be interpreted as beliefs, but rather incorporates a much wider set of positive feelings about an action, as well as notions of forgetting or confirmatory bias. This is why this rule is typically given the interpretation of a reinforcement or stimulus response rule.¹⁴ Note that under either rule experience of payoffs from action m impacts propensities in all games belonging to the same category g but not in games outside category g .

We are not so much interested now in distinguishing between these two learning models, but rather we want to use them as a workhorse for analyzing which categorizations explain behavior best. Our estimation strategy is the following. We will estimate the probability of choosing a particular action in a particular game under both learning models and under all possible assumptions on the categories used by the agents. We will then evaluate model accuracy (of all these different models) by using the Akaike information criterion. Since we are interested in whether different action choices of agents are due to different categories the agents hold, we estimate the probability of choosing a given action, i.e.

$$\Pr(a(t) = m|g)^{it} = \frac{\exp \theta \sum_{h \in \{A,B,C\}} \beta_h \alpha_h^{it}(g)}{1 + \exp \theta \sum_{h \in \{A,B,C\}} \beta_h \alpha_h^{it}(g)}.$$

We then check whether different categories explain behavior best for different action choices. We restrict attention to the last 50 periods of the experiment, where behavior has stabilized. The regression tables for the "best" regressions are provided in the appendix.

Our model has several free parameters. We assume four different values for $\delta \in \{0.05, 0.2, 0.5, 0.8\}$ and different initial conditions for $(\alpha_A^{i0}, \alpha_B^{i0}, \alpha_C^{i0})$. In particular we set initial conditions to either the zero vector or proportional to first round choice (summing to 40). We then estimate all possible combinations of these values and the different categories and learning models and use the Akaike information criterium to select among all models. Our main findings are summarized in the following tables, where we restrict attention to row players, since column players have a dominant

"Best" Categories

The table illustrates the "best partition" (according to the Akaike criterium) for each of the action choices and games. It also shows which model (the belief-based model or the reinforcement model) explained the data better, and which value of δ . The "best"

¹⁴In again another class of models information about counterfactual payoffs are used to update propensities. We do not include such models here, since in treatment M using such information essentially means memorizing the entire payoff matrix.

	Game1	Game4	Game3	Game6
<i>A</i>	– (0.42)	{1, 2, 3, 4, 5, 6} BB, $\delta = 0.2$ (0.48)	{3, 6} RF, $\delta = 0.05$ (0.54)	{3, 6} BB, $\delta = 0.05$ (0.50)
<i>B</i>	{1, 3, 4, 6} BB, $\delta = 0.2$ (0.46)	{1, 3, 4, 6} BB, $\delta = 0.2$ (0.46)	– (0.00)	– (0.00)
<i>C</i>	– (0.09)	– (0.06)	{1, 3, 4, 6} RF, $\delta = 0.05$ (0.46)	{1, 2, 3, 4, 5, 6} RF, $\delta = 0.05$ (0.50)

Table 8: "Best" partitions for each action choice and game, "best" model according to the AIC (BB=Belief Based, RF=Reinforcement), and shares of the corresponding choices (in brackets).

δ was usually either 0.05 or 0.2. Initial conditions did not make any substantial difference in terms of the Akaike criterium. (Remember also that we focus on the last 50 periods). For some regressions (e.g. A-choices in game 1) no significant learning can be singled out by our regressions. Some regressions did not produce enough variation to yield significant results. This is especially true for choices that were observed very little.

First, observe that strategically equivalent games are always found in the same category. The most clear-cut results are obtained in games 3 and 6. Remember that in these games column players almost always choose strategy b and learning thus might be "easier". Nash choices (*A*) are always best explained through category {3, 6}, whereas Non-Nash choices (*C*) are best explained through the bigger categories {1, 3, 4, 6} or even {1, 2, 3, 4, 5, 6}. In fact, for games 3 and 6 these two categories are very close in terms of the Akaike information criterium. As we stated in section 2.2 (and prove in the appendix), a model of *action bundling* predicts choosing *C* in these games for category {1, 2, 3, 4, 5, 6}. Belief bundling cannot explain *C*-choices in these games. Also the fact that reinforcement learning explains the data best is consistent with the hypothesis that action bundling led to the observed *C*-choices. *A*-choices in category {3, 6} are predicted by both models (action and belief-bundling).

In games 1 and 4 *B*-choices are best explained by category {1, 3, 4, 6}. *Belief bundling* predicts *B*-choices given category {1, 3, 4, 6}, but also action bundling predicts *B*-choices if historical frequencies are used (see again section 2.2 and the appendix). The only somewhat puzzling result is that *A*-choices in Game 4 seem to be explained through bigger categories than *B*-choices. We will come back to this below.

Note also that table 8 does not answer our research question completely. We would like to know whether our results actually reflect the fact that there are different "player types", each holding consistent categories. Thus, we repeat the previous exercise, but

now focusing on agents that choose action A , B or C the majority of the time. Table 9 reports the results.

	Game1	Game4	Game3	Game6
A	– (0.48)	{4} BB, $\delta = 0.2$ (0.44)	{3, 6} RF, $\delta = 0.05$ (0.50)	{3, 6} BB, $\delta = 0.05$ (0.50)
B	– (0.46)	{1, 3, 4, 6} BB, $\delta = 0.2$ (0.50)	– (0.00)	– (0.00)
C	– (0.06)	– (0.06)	{1, 3, 4, 6} RF, $\delta = 0.05$ (0.50)	{1, 2, 3, 4, 5, 6} RF, $\delta = 0.05$ (0.50)

Table 9: For agents that choose mainly A (B , C): "Best" partitions for each action choice and game, "best" model according to the AIC (BB=Belief Based, RF=Reinforcement), and shares of the corresponding choices (in brackets).

Now the puzzling result in Game 4 disappears. Among those agents that mainly choose action A the finest category $\{4\}$ indeed explains the data best. All other results stay essentially unchanged, which makes us confident that indeed our results observed in Section 3.1 can be explained well by models of categorization. Lastly the fact that the best δ is somewhat lower in games 3 and 6 could indicate that in these games behavior has converged faster than in the other games.

In order to know whether similarity relations evolve over time we repeated the exercise by focusing on the first 33 rounds of the experiment. Here we find that (in almost all cases) the coarsest category explains the data best, suggesting that categories are indeed "learned" over time.

RESULT 4 (CATEGORIZATION AND ACTION CHOICES) *1. Nash choices (A in games 1,3,4 and 6) are best explained through smaller categories than non-Nash choices.*

2. B choices in games 1 and 4 can be rationalized through models of belief bundling or action bundling given the "best" category.

3. C choices in games 3 and 6 can be rationalized through models of action bundling given the "best" categories.

Population Heterogeneity

While our results can be explained well by models of categorization with either action or belief-bundling, still a somewhat puzzling feature remains. After all we would like to think that agents hold a number of categories and choose consistently with

this categorization (and a model of either belief-bundling or action bundling) in all games. No problem arises in this respect with A -choices. But with respect to the Non-Nash choices the picture is a little less clear. For example, if an agent holds category $\{1, 3, 4, 6\}$ and engages in action bundling she should choose the same action (either A or B) in all games 1, 3, 4 and 6. On the other hand if an agent engages in belief-bundling she should choose action B in games 1 and 4 and action A in games 3 and 6, but never action C . In trying to reconcile these facts we will also touch our question of transitivity of categorizations.

Non-Transitivity of Categorizations

A first explanation would be that agents choosing Non-Nash actions use the biggest category for their decisions in games 3 and 6, i.e. extrapolate from all games in games 3 and 6 but hold the smaller category $\{1, 3, 4, 6\}$ for decisions in games 1 and 4. If we furthermore assume that agents engage in action bundling this could explain the results, but we would have to accept that categorizations are not partitions and in particular are neither transitive nor symmetric. A more convincing explanation is maybe to maintain the partition structure of categorizations and rely exclusively on heterogeneity in the population to explain the results.

Population Heterogeneity w.r.t. categories

The results can also be explained if we assume that roughly 50% of both row and column players hold the biggest category and 50% hold partition $\{\{1, 4\}, \{2, 5\}, \{3, 6\}\}$. Assume also that all players engage in action bundling. Column players holding the biggest category should then choose action b in all games (since it is dominant strategy given this category). Row players holding this category should choose either action C or action A (see appendix). To explain our results we in fact need that agents resolve this indifference much more often in favor of action C in games 3 and 6 than in games 1 and 4. This (together with the fact that belief-based learning shows up in the best regressions) could indicate that if not belief-bundling at least some belief-based reasoning is also present in the data. Note that this explanation could account for the result in Table 8 where A -choices are explained through a bigger category than B -choices in game 4.

Column players holding categorization $\{\{1, 4\}, \{2, 5\}, \{3, 6\}\}$ should choose a in $\{1, 4\}$ and b in the remaining categories. Row players on the other hand holding categorization $\{\{1, 4\}, \{2, 5\}, \{3, 6\}\}$ should choose A in $\{3, 6\}$, but might choose either A or B in category $\{1, 4\}$, since both actions are best responses to column players choosing $\frac{1}{2}a \oplus \frac{1}{2}b$.

An alternative explanation could also be to split the latter group into two groups where one holds categorization $\{\{1, 4\}, \{2, 5\}, \{3, 6\}\}$ and one holds $\{\{1, 3, 4, 6\}, \{2, 5\}\}$. Note also that the heterogenous population explanation is consistent with the evidence presented in Section 3.3.

Population Heterogeneity w.r.t. analogy models and categories

Another possibility is allowing a heterogenous population along both dimensions. If we make the same assumptions as above but now assume in addition that all agents holding categorization $\{\{1, 4\}, \{2, 5\}, \{3, 6\}\}$ engage in belief-bundling, the results can be equally well explained. Note also that the previous paragraphs of course illustrate only the minimal amount of heterogeneity needed to explain the results. Unfortunately the data are too noisy to get results on population heterogeneity which are of a less speculative nature.

3.3 Cognitive Reflection Test

Since the aim of the paper was to investigate decisions in the presence of cognitive constraints, it is natural to expect that the willingness of subjects to engage in cognitively costly reasoning processes should be correlated with their behavior in the different games. To these ends we conducted a cognitive reflection test at the end of the experiment (Fredericks, 2005). The test consists of the three following questions

1. A bat and a ball cost Euro 1.10 in total. The bat costs Euro 1.00 more than the ball. How much does the ball cost ?
2. If it takes 5 machines 5 minutes to make 5 widgets, how long would it take 100 machines to make 100 widgets ?
3. In a lake, there is a patch of lily pads. Every day, the patch doubles in size. If it takes 48 days for the patch to cover the entire lake, how long would it take for the patch to cover half of the lake ?

All these questions have an answer that immediately springs into mind (10 cents, 100 minutes, 24 days), which is wrong. The right answer (5 cents, 5min, 47 days) can only be found by engaging in some cognitive reflection. Note that the test is not measuring intelligence, but rather the willingness of subjects to engage in costly cognitive reflection. Table 11 summarizes the results from the cognitive reflection test in treatment M .

Question No.	1	2	3
Reflective	47%	41%	47%
Unreflective	50%	50%	34%
Other Answers	3%	9%	19%

Table 10: Reflective, unreflective and other answers in the CRT in treatment M (percentage of individuals).

We now report some results on the relation of the individuals' CRT scores and their choice behavior in our experiment. In the next table we report the action choices

in games 1,3,4 and 6 in treatment M separately for those subjects who got all three questions right and for those that did not. Getting all three questions right implies a very high willingness to engage in costly cognitive reflection. Hence we will refer to these subjects as "very reflective".

Game 1	A	B	C
very reflective	40.7%	57.1%	2.2%
others	44.6%	42.2%	13.2%
Game 4	A	B	C
very reflective	39.5%	58.1%	2.3%
others	37%	48.5%	14.1%
Game 3	A	B	C
very reflective	67.6%	–	32.3%
others	37%	1.5%	59.6%
Game 6	A	B	C
very reflective	87.5%	–	12.5%
others	38.1%	1.2%	60.7%

Table 11: Action choice of reflective and unreflective individuals in treatment M .

Focus first on games 3 and 6. Here it can be seen that very reflective subjects choose C a lot less than the best response action A . Their higher willingness to engage in costly cognitive reflection seems to make them willing to use finer categories than other agents. It is also interesting that the reflective subjects choose the suboptimal action C (which yields a lot lower payoff given the choices of others) a lot less in games 1 and 4. There seem to be no significant differences between the proportions of A and B choices in these games. If anything the very reflective agents tend to use action B more often in those games. This is consistent with the heterogenous population explanation given above. Note also that the results are again extremely similar in games 1 and 4 and games 3 and 6 respectively. We can summarize as follows.

RESULT 5 (COGNITIVE REFLECTION TEST) *Row Players that are most willing to engage in costly cognitive reflection choose Action C significantly less often in all games 1, 3, 4, and 6 than other row players.*

4 Conclusions

In this paper we investigated experimentally how agents learn to make decisions in a multiple games environment. Participants interacted in simple 3×3 normal form games for 100 rounds. We varied the number of games across treatments. In addition in some treatments we provided explicit information about the opponent's past behavior. We find that participants do extrapolate between games. If either there

are few games or if explicit information about the opponent's behavior is provided (or both) convergence to the unique Nash equilibrium generally occurs. Otherwise this is not the case and play converges to a distribution of actions which is Non-Nash. Action choices, though, that cannot be explained by theoretical models of either belief-bundling or action bundling are never observed. Estimating both belief-based and reinforcement learning models using different partitions we find that Nash choices are best explained by finer categorizations than Non-Nash choices. Furthermore participants scoring better in the "Cognitive Reflection Test" choose Nash actions more often than other participants.

References

- [1] Abreu, D. and A. Rubinstein (1988), The Structure of Nash Equilibrium in Repeated Games with Finite Automata, *Econometrica* 56(6), 1259-1281.
- [2] Akaike, H. (1976), An information criterion, *Math.Sci.*14: 5-9.
- [3] Burnham, K.P., and D.R. Anderson (2002), *Model Selection and Multimodel Inference: A Practical-Theoretic Approach*, 2nd Edition. Springer, New York.
- [4] Camerer, C. and T. Ho (1999), Experience-weighted attraction learning in normal form games, *Econometrica* 67, 827-874.
- [5] Frederick, S. (2005), Cognitive Reflection and Decision Making, *Journal of Economic Perspectives* 19(4), 25-42.
- [6] Gilboa, I. and D. Schmeidler (1995), Case-Based Decision Theory, *The Quarterly Journal of Economics*, 110(3), 605-639.
- [7] Gilboa, I. and D. Schmeidler (1996), Case-Based Optimization, *Games and Economic Behavior* 15, 1-26.
- [8] Jehiel, P. (2005), Analogy-based expectation equilibrium, *Journal of Economic Theory* 123, 81-104.
- [9] Haruvy, E. and D.O. Stahl (2008), Learning Transference between Dissimilar Symmetric Normal Form Games, mimeo, University of Texas.
- [10] Hopkins, E. (2002), Two Competing Models of How People Learn in Games, *Econometrica* 70(6), 2141-2166.
- [11] Hopkins, E. (2007), Adaptive Learning Models Of Consumer Behavior, *Journal of Economic Behavior and Organization* 64, 348-368.
- [12] Huck, S., P. Jehiel and T. Rutter (2007), Learning Spillover and Analogy Based Expectations: A Multi Game Experiment, working paper UCL.

- [13] Mengel, F. (2008), Learning Across Games, working paper Maastricht University.
- [14] Roth, A.E. and I. Erev (1995), Learning in Extensive-Form Games: Experimental Data and Simple Dynamic Models in the Intermediate Term, *Games and Economic Behavior* 8, 164-212.
- [15] Rubinstein, A. (1988), Similarity and Decision-making under Risk (Is There a Utility Theory Resolution to the Alais Paradox?), *Journal of Economic Theory* 46, 145-153.
- [16] Rustichini, A. (1999), Optimal Properties of Stimulus Response Learning Models, *Games and Economic Behavior* 29, 244-273.
- [17] Samuelson, L. (2001), Analogies, Anomalies and Adaptation, *Journal of Economic Theory* 97, 320-366.
- [18] Sarin, R. and Vahid (1999), Payoff Assessments without Probabilities: A Dynamic Model of Choice, *Games and Economic Behavior* 28, 294-309.
- [19] Stahl, D.O. and J. van Huyck (2002), Learning Conditional Behavior in Similar Stag Hunt Games, mimeo, University of Texas.
- [20] Steiner, J. and C. Stewart (2007), Learning by Similarity in Coordination Problems, mimeo CERGE-EI.
- [21] Young, P. (1993), The Evolution of Conventions, *Econometrica* 61(1), 57-84.

5 Appendix

5.1 Predictions Belief Bundling and Action Bundling

We will consider categories $\{1, 3, 4, 6\}$ as well as $\{1, 2, 3, 4, 5, 6\}$ to illustrate how action bundling and belief-bundling do work.

Action Bundling

Consider the average game created through category $\{1, 2, 3, 4, 5, 6\}$. Denote by f_i the frequency of game i . Then payoffs for the row player in the average game can be computed as follows:

	a	b	c
A	$20(f_1 + f_6) + 5f_2 + 15f_3 + 25f_4 + 15f_5$	$15f_1 + 20(f_3 + f_4) + \dots$	$15(f_1 + f_3) + \dots$
B	$10f_1 + 15(f_3 + f_4) + 20f_6 \dots$	$25f_1 + 10f_3 + \dots$	$5(f_3 + f_4) + \dots$
C	$10f_1 + 20f_3 + 15f_4 + 25f_6 \dots$	$15(f_1 + f_3) + \dots$	$35(f_1 + f_3) + \dots$

If we now apply uniform frequencies and add the payoffs for column players we get the average game

	a	b	c
A	16.7, 15	20, 20	15, 15
B	15, 20	16.7, 20	8.3, 13.3
C	20, 12.5	21.7, 31.7	31.7, 6.7

If agents engage in action bundling and play according to the average game we should observe an increase in C choices for row players and an increase of b -choices for column players, since both are dominant strategies given this average game. This continues to be true if historical frequencies from the first 50 rounds are used instead of uniform frequencies, but in this case action C ceases to be a dominant strategy for row players (also A choices are rationalizable), while b is still the dominant strategy for column players.

For category $\{1, 3, 4, 6\}$ we get

	a	b	c
A	20, 17, 5	20, 17, 5	17.5, 15
B	15, 20	20, 15	5, 12, 5
C	17.5, 13.75	17.5, 37, 5	37.5, 2, 5

It can be seen easily that action c is dominated and should thus never be observed. The predicted outcome is an increase in b -choices by column players compared to the Nash benchmark.¹⁵ If historical frequencies $(f_1, f_3, f_4, f_6) = (0.21, 0.24, 0.48, 0.06)$ are used instead we find (for row players)

	a	b	c
A	21.9	19.05	17.55
B	14.1	22.95	4.2
C	15.6	17.55	37.35

Using these frequencies thus action bundling predicts an increase also in B -choices by row players whenever at least $\frac{2}{3}$ of column players choose b .¹⁶

Belief Bundling

Under belief bundling agents form average beliefs but then best respond to these average beliefs in each game separately.

In category $\{1, 2, 3, 4, 5, 6\}$, since choosing b is a dominant strategy in games 2 and 5 and results from iterated elimination of dominated strategies in games 3 and 6, row players must attach belief at least $\frac{2}{3}$ to the column player choosing b (and probability

¹⁵Note that these results depend in principle on the frequencies used in the calculations. If we use actual frequencies (ex post) instead the results though stay unchanged.

¹⁶If we focus on average behavior in rounds 25 – 50, then action bundling predicts an increase in B -choices whenever at least 40% of column players choose action b .

zero to c). But then B is a best response to any such belief in games 1 and 4. Given that row players choose B in these games (i.e. with frequency $\frac{1}{3}$), C in games 2 and 5 (frequency $\frac{1}{3}$) and A in games 3 and 6 (frequency $\frac{1}{3}$) choosing b by the column players is a best response in turn. Consequently belief-bundling predicts an increase in B (b) choices in games 1 and 4 under this category.

In category $\{1, 3, 4, 6\}$ again b is a dominant strategy for column players in games 3 and 6 (after eliminating B) and consequently must receive at least probability $\frac{1}{2}$. But then again B is a best response for row players in games 1 and 4. Given that row players choose B with probability $\frac{1}{2}$ and A with probability $\frac{1}{2}$ (games 3 and 6) column players will best respond with A in games 1 and 4 leading. Belief-bundling predicts an increase in B choices in games 1 and 4 in these games.

5.2 Regression Tables: Best Partitions in Games 3,4 and 6.

A	Coeff.	Std.Error	$P > z $	95% Interval
constant	1.500011	1.2385	0.226	[-0.92,2.40]
$\alpha_A^{BB}(\{3, 6\})$	2.053885	0.99659	0.040	[0.09,3.22]
$\alpha_B^{BB}(\{3, 6\})$	0.122238	0.16964	0.471	[-0.21,0.34]
$\alpha_C^{BB}(\{3, 6\})$	-0.646105	0.23883	0.007	[-1.11,-0.13]
σ_u	0.055531			
ρ	0.093000			

A-choices in Game 3, category $\{3, 6\}$

A	Coeff.	Std.Error	$P > z $	95% Interval
constant	-1.393765	1.19398	0.243	[-3.73,1.66]
$\alpha_A^{RF}(\{3, 6\})$	0.6715005	0.1885119	0.000	[0.30,0.93]
$\alpha_B^{RF}(\{3, 6\})$	0.1140251	0.9464411	0.904	[-1.74,1.85]
$\alpha_C^{RF}(\{3, 6\})$	-0.1069534	0.0218569	0.062	[-0.53,0.33]
σ_u	1.320635			
ρ	0.3464633			

A-choices in Game 6, category $\{3, 6\}$

C	Coeff.	Std.Error	$P > z $	95% Interval
constant	2.791228	0.695832	0.000	[1.42,3.81]
$\alpha_A^{BB}(\{1, 3, 4, 6\})$	-0.217527	0.0691769	0.002	[-0.35,-0.07]
$\alpha_B^{BB}(\{1, 3, 4, 6\})$	-0.3996449	0.0747192	0.000	[-0.54,-0.26]
$\alpha_C^{BB}(\{1, 3, 4, 6\})$	0.1142352	0.0169873	0.000	[0.08,0.19]
σ_u	1.80657			
ρ	0.4650672			

C-choices in Game 3, category $\{1, 3, 4, 6\}$

C	Coeff.	Std.Error	$P > z $	95% Interval
constant	1.272907	1.82598	0.486	[-2.30,4.01]
$\alpha_A^{RF}(\{1, 2, 3, 4, 5, 6\})$	-0.2361806	0.1479621	0.110	[-0.52,-0.09]
$\alpha_B^{RF}(\{1, 2, 3, 4, 5, 6\})$	-0.2481736	0.1229855	0.044	[-0.48,-0.10]
$\alpha_C^{RF}(\{1, 2, 3, 4, 5, 6\})$	0.2374154	0.0895438	0.008	[0.06,0.41]
σ_u	1.320635			
ρ	0.3464633			

C-choices in Game 6, category $\{1, 2, 3, 4, 5, 6\}$

A	Coeff.	Std.Error	$P > z $	95% Interval
constant	1.925352	3.19106	0.546	[-4,32,6.20]
$\alpha_A^{BB}(\{1, 2, 3, 4, 5, 6\})$	0.3551018	0.1155269	0.002	[0.12,0.59]
$\alpha_B^{BB}(\{1, 2, 3, 4, 5, 6\})$	-0.714897	0.1606167	0.000	[-1.02,-0.39]
$\alpha_C^{BB}(\{1, 2, 3, 4, 5, 6\})$	0.1597932	0.0777148	0.040	[0.00,0.29]
σ_u	1.089347			
ρ	0.2650876			

A-choices in Game 4, category $\{1, 2, 3, 4, 5, 6\}$

B	Coeff.	Std.Error	$P > z $	95% Interval
constant	-0.57705905	4.668256	0.902	[-2.30,4.01]
$\alpha_A^{BB}(\{1, 3, 4, 6\})$	-0.3709647	0.161549	0.022	[-0.68,-0.07]
$\alpha_B^{BB}(\{1, 3, 4, 6\})$	0.4667594	0.157375	0.003	[0.15,0.61]
$\alpha_C^{BB}(\{1, 3, 4, 6\})$	-0.1126542	0.092494	0.226	[-0.29,0.11]
σ_u	1.648653			
ρ	0.452412			

B-choices in Game 4, category $\{1, 3, 4, 6\}$