

Fehr, Dietmar; Kübler, Dorothea; Danz, David

Working Paper

Information and beliefs in a repeated normal-form game

WZB Discussion Paper, No. SP II 2010-02

Provided in Cooperation with:
WZB Berlin Social Science Center

Suggested Citation: Fehr, Dietmar; Kübler, Dorothea; Danz, David (2010) : Information and beliefs in a repeated normal-form game, WZB Discussion Paper, No. SP II 2010-02, Wissenschaftszentrum Berlin für Sozialforschung (WZB), Berlin

This Version is available at:
<http://hdl.handle.net/10419/54567>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

WZB

WISSENSCHAFTSZENTRUM BERLIN
FÜR SOZIALFORSCHUNG

SOCIAL SCIENCE RESEARCH
CENTER BERLIN

Dietmar Fehr *
Dorothea Kübler **
David Danz *

**Information and Beliefs in a
Repeated Normal-form Game**

* WZB

** WZB, TU Berlin & IZA

SP II 2010 – 02

March 2010

ISSN Nr. 0722 – 6748

**Research Area
Markets and Politics**

**Research Unit
Market Behavior**

**Schwerpunkt II
Märkte und Politik**

**Abteilung
Verhalten auf Märkten**

Zitierweise/Citation:

Dietmar Fehr, Dorothea Kübler, David Danz, **Information and Beliefs in a Repeated Normal-form Game**, Discussion Paper SP II 2010 – 02, Wissenschaftszentrum Berlin, 2010.

Wissenschaftszentrum Berlin für Sozialforschung gGmbH,
Reichpietschufer 50, 10785 Berlin, Germany, Tel. (030) 2 54 91 – 0
Internet: www.wzb.eu

ABSTRACT

Information and Beliefs in a Repeated Normal-form Game*

by Dietmar Fehr, Dorothea Kübler, and David Danz

We study beliefs and choices in a repeated normal-form game. In addition to a baseline treatment with common knowledge of the game structure, feedback about choices in the previous period and random matching, we run treatments (i) with fixed matching, (ii) without information about the opponent's payoffs, and (iii) without feedback about previous play. Using Stahl and Wilson's (1995) model of limited strategic reasoning, we classify behavior with regard to its strategic sophistication and consider its development over time. In the treatments with feedback and full information about the game, we observe more strategic play, more best-responses to beliefs and more accurate beliefs over time. While feedback is the main driving force of learning to play strategically and for forming beliefs that accurately predict the behavior of the opponent, both incomplete information about the opponent's payoffs or lack of feedback lead to a stagnation of best-response rates over time.

Keywords: experiments, beliefs, strategic uncertainty, learning

JEL Classification: C72, C92, D84

*

For valuable comments, we thank seminar participants at the Humboldt-Universität zu Berlin, European University Florence, SFB 649 Workshop 2007, ESA World Meeting 2007, IMEBE 2008, VfS Annual Meeting 2008 and Econometric Society Meetings 2009 as well as Kyle Hyndman, Harald Uhlig, Roberto Weber, Georg Weizsäcker and Axel Werwatz. We are indebted to Jana Stöver and Susanne Thiel for research assistance. Financial support from the Deutsche Forschungsgemeinschaft (DFG) through the SFB 649 "Economic Risk" is gratefully acknowledged. Corresponding Author: Dorothea Kübler, Social Science Research Center Berlin (WZB), Research Unit "Market Behavior", Reichpietschufer 50, 10785 Berlin. Email: kuebler@wzb.eu.

ZUSAMMENFASSUNG

Information und Erwartungen in einem wiederholten Normalformspiel

Wir untersuchen die Entwicklung von den Erwartungen über das Verhalten des anderen Spielers und den Entscheidungen in einem wiederholten Normalformspiel. Zusätzlich zum Haupttreatment mit common knowledge über das Spiel, Feedback über das Ergebnis in der vorigen Runde und zufälliger Zuordnung der Spieler, gibt es Kontrolltreatments mit (i) festen paarweisen Zuordnungen der Spieler, (ii) ohne Information über die Auszahlungen des anderen Spielers und (iii) ohne Feedback über das Ergebnis der vorigen Runde. Mit Hilfe von Stahl und Wilsons (1995) Modell begrenzten strategischen Verhaltens klassifizieren wir das Verhalten der Teilnehmer im Hinblick auf die strategische Sophistikation. In den Treatments mit Feedback und vollständiger Information über das Spiel nehmen strategisches Verhalten, beste Antworten auf die eigenen Erwartungen und die Akkuratheit der Erwartungen über die Zeit zu. Während Feedback der Hauptgrund dafür ist, dass die Teilnehmer lernen, sich strategisch zu verhalten und korrekte Erwartungen über das Verhalten des anderen Spielers zu bilden, führen sowohl unvollständige Information über die Auszahlungen des Gegenspielers als auch fehlendes Feedback zu einer Stagnation der Rate der besten Antworten über die Zeit.

1 Introduction

The literature on learning has opened the black box of how an equilibrium is reached. Numerous theoretical and experimental papers have studied learning over a large number of periods and have focused either on the convergence properties of the learning algorithms or on the evolution of observed behavior in experimental data. Most learning models are backward looking and model decisions using past observations. More sophisticated learning models posit a deductive reasoning process implying that players analyze the game in order to understand its strategic properties and thereby form beliefs about the opponent's choice. In this paper we take a microscopic view of the learning process in order to disentangle its inductive and deductive elements. By varying the information conditions, we control for the impact of experience and sophistication. Thus, we provide a unified framework to study deductive learning in a no-feedback environment and experience-based inductive learning in an environment where relevant information for forward-looking learning is lacking.

For the experiment, we use a normal-form game with a unique Nash equilibrium that is Pareto-dominated. Beliefs are measured using an incentive compatible elicitation procedure. Thus, we complement the decisions with subjects' elicited beliefs and observe the joint development of beliefs and decisions over time. The game we chose allows for a clear-cut distinction between strategies with higher and lower levels of strategic thinking in the sense of Stahl and Wilson's level-k model (1995). Using this classification of strategies, we can track the change in the level of strategic choices of players over time. In the game we use, the Nash equilibrium is Pareto-dominated by another outcome. Therefore, the game allows for a differentiation between strategic types with purely self-interested preferences and strategic types with other-regarding preferences, and it requires players to form beliefs about other players' types.

First, we run a baseline treatment with full information about the game and with feedback about one's own payoff (and thereby the other's payoff and action) in the previous period. In this treatment, we use a random matching protocol. To check whether the observed learning patterns are robust to changes in the matching procedure, we employ a treatment with fixed pairs for the whole experiment of 20 periods. These two treatments with full information about the game and past outcomes serve as our main treatments. To be able to separate between the different forms of learning, we employ two additional control treatments. To account for the possibility of sophistication without feedback (see e.g. Weber 2003), we use a treatment in which subjects receive

no feedback about the current play.¹ In the second control treatment, subjects know only their own payoff function and receive feedback about previous play, but they do not know the payoff function of the other player.² Studying learning in a normal-form game under both information conditions allows us to compare their relative importance for learning. Note that sophisticated learners use the information about the other player's payoffs which is not used by purely experience-based learners while the experienced-based learners make use of feedback information.

We find an initially high level of non-strategic behavior in all treatments as subjects tend to neglect the incentives of their opponents. In the three treatments with feedback about the other player's past behavior, this non-strategic behavior decreases over time and Nash play increases. In the treatment without feedback about past outcomes, there is virtually no change in behavior over time. Thus, our results indicate the importance of feedback. Information about the other player's payoffs matters much less in that it is important for initial play, but much less than expected from rational players. Thus, subjects seem to have only a limited understanding of the strategic properties of the game initially, even when they have full information about the game. Also, the development of choices over time is very similar in treatments with and without information about the opponent's payoff function.

Regarding the beliefs, we first confirm that stated beliefs are better predictors of the actual choices than the estimated beliefs using belief-learning models. Therefore, we work with the stated beliefs in all subsequent analyses. Both in standard Nash equilibrium and in the level-k model, players are assumed to best respond to their beliefs. However, best-response rates are initially only between 50% and 60% in all treatments. We observe an increase in best responses over time in the two treatments with full information about the game and about past outcomes, but not in the two control treatments. Thus, information about past play of the opponent and about his incentives in the game allow subjects to learn to best respond. As both the Nash equilibrium concept and the level-k model do not allow for such failures to best respond, this form of learning is not captured by them.

¹It is conceivable that experience and observation of past play could reduce the need of sophistication. In a feedback-free environment subjects are presumably more forced to think about the game and therefore they may acquire simple solution concepts such as iterated dominance or backwards induction. Weber and Rick (2008) demonstrate that subjects are able to acquire and to transfer such concepts to similar games, but only in feedback-free environments.

²Oechssler and Schipper (2006) used a similar setup to study subjects' ability to learn about the game they are playing.

In the framework of the level-k model, players may hold inaccurate beliefs because some types assume that they are more sophisticated than other players. But players may in fact become more sophisticated and increase the number of steps of reasoning in the course of the experiment such that they form more accurate beliefs later on. Thus, the belief data provide evidence on the level of reasoning of the subjects. Again, we find that information about the opponent's past choices is necessary for improving the accuracy of belief statements over time.

There are a few recent papers using a belief elicitation procedure in finitely repeated normal-form games. Nyarko and Schotter (2002) investigated the explanatory power of beliefs inferred from belief-learning models such as fictitious play models. They used a 2x2 game with a unique mixed-strategy equilibrium and found that belief learning models cannot predict stated beliefs well. The two closest papers to our design are Ehrblatt, Hyndman, Özbay and Schotter (2008) and Terracol and Vaksman (2009). The first paper focuses on strategic teaching and its underlying mechanisms using two normal-form games with a unique Nash equilibrium that is Pareto efficient. The authors demonstrate that the convergence process largely depends on the presence of a sophisticated subject, the teacher, and a fast enough follower. Terracol and Vaksman (2009) also investigate learning and teaching, but in a game with multiple non-Pareto rankable equilibria. They find evidence for self-interested teaching, but the multiplicity of equilibria creates a conflict between the players, resulting in a slower convergence process. All three studies have in common that they do not focus on the relative importance of deductive and inductive learning for the evolution of strategic play in a game. Although Ehrblatt et al. (2008) also ran a treatment with incomplete information about the opponent's payoffs, none of the papers employs a treatment without feedback. Furthermore, the Pareto-dominated Nash equilibrium in our game allows for a differentiation of strategic types with respect to their social preferences, making the belief formation task more demanding.

The paper is organized as follows. The next section introduces the design and procedures of the experiment and provides a description of the level-k model applied to the normal-form game we used. In Section 3, we present the results, focusing first on choices and then on belief statements. Section 4 concludes.

	Left	Center	Right
Top	78, 68	72, 23	12, 20
Middle	67, 52	59, 63	78, 49
Bottom	21, 11	62, 89	89, 78

Table 1: Game.

2 Experimental design

2.1 The game

In all treatments of the experiment, we used the asymmetric normal-form game presented in Table 1. The game has a unique Nash equilibrium in pure strategies in which the row player chooses Top and the column player chooses Left. This equilibrium can be found by applying iterative elimination of dominated strategies. Note that the Nash equilibrium of the stage game is not Pareto efficient. The strategy combination of Bottom and Right leads to higher payoffs for both players. This outcome maximizes the payoff of the player that is least well off, and it also maximizes the sum of payoffs. As we are interested in the relationship between beliefs and choices, we chose a game where beliefs about the other player’s preferences can affect behavior.

The unique Nash equilibrium of the stage game is also the unique subgame perfect equilibrium of the repeated game. However, there exist Nash equilibria of the finitely repeated game with fixed matching in which the players choose the Pareto-efficient strategy combination (Bottom, Right) for a number of periods and then switch to the Nash Equilibrium (Top, Left).³ Finally, note that for the column player choosing Right is strictly dominated by Center.

2.2 Strategies

Stahl and Wilson (1995) proposed a theory of boundedly rational types, based on a hierarchical model by Nagel (1993). Stahl and Wilson assume that players differ in their level of strategic sophistication. Their model classifies players into types according to their level of reasoning. A level-0 type randomizes uniformly over his strategy space, whereas a level- k type best responds to level- $(k - 1)$ behavior for $k \in \{1, 2, \dots, \infty\}$. Hence the term level- k model.⁴

³In case a player deviates in this equilibrium, she is minmaxed by the other player choosing Middle or Center, respectively, for the rest of the game.

⁴The level- k model is a useful approach to track off-equilibrium behavior. It has been tested and extended by various other studies mainly in the context of normal-form games (e.g. Costa-Gomes et al 2001; Costa-Gomes and

	Row player		Column player	
Top	L2+/Nash	Left	L3+/Nash	
Middle	L1	Center	L1/L2	
Bottom	Utilitarian	Right	Utilitarian	

Table 2: Decision rules.

The level-k model is a static model, but in our repeated setting learning becomes possible. Within the level-k model, learning can be understood as subjects choosing higher-level strategies. Suppose a subject starts out by playing the $L1$ action, but then learns to best respond to $L1$ by playing $L2$ and so forth. Thus, a subject can learn by updating his beliefs in the course of the game, and we will investigate this on the basis of our data. In particular, we will test whether the subjects choose higher-level strategies and whether beliefs become more accurate in predicting the opponents' behavior over time.

The main focus of this study is on the development of strategic and non-strategic behavior over time. Thus we use the level-k model to classify the available strategies in our game (see Table 2) and distinguish between strategic and non-strategic types. Strategic types form beliefs based on an analysis of what others do and best respond to these beliefs, whereas non-strategic types do not take into account the incentives of others. Given this definition, a strategic, self-interested row player would choose Top ($L2+$) and a strategic column player Left ($L3+$) or Center ($L2$). It also emerges from Table 2 that two and three steps of thinking, respectively, are sufficient to reach the unique Nash equilibrium of the stage game.

As the Nash equilibrium is not Pareto efficient, we can distinguish between Nash play and play of the most efficient and/or fair outcome. In our game, it is possible that subjects play higher-level strategies in order to maximize joint payoffs. This behavior cannot be identified in games where the Nash equilibrium is on the Pareto frontier. Thus, we also introduce a joint-payoff maximizing (or Utilitarian) decision rule, which maximizes the sum of the payoffs of both players, given that the other player has the same objective and chooses accordingly.⁵ According to the Weizsäcker 2008; Rey Biel forthcoming; Ivanov 2006; or Camerer et al 2004). It is also successful in organizing data from other games such as auctions, as recently shown by Crawford and Iriberry (2007a, 2007b) as well as Gneezy (2005). The most common types found in normal-form games are level-1 ($L1$), level-2 ($L2$) and Nash types, but their distribution crucially depends on the set of games investigated.

⁵Previous studies did not explicitly explore Utilitarian choices, but some of them found behavior pointing in this direction (e.g. Costa-Gomes and Weizsäcker, 2008).

proposed definition of strategic behavior, the joint-payoff maximizing action is strategic because it requires the belief that the other player has the same preferences and acts accordingly. Hence, under the assumption of other-regarding preferences, the strategies Middle and Center correspond to L1 whereas all higher-level strategies coincide with Bottom and Right.

The game chosen allows us to identify strategic and non-strategic behavior as clearly as possible, while we can also distinguish the players with respect to their preferences. For the sake of a simple classification of the actions in Table 2 in terms of their strategic sophistication, we proceed as follows. Since the actions Middle and Center represent best responses to random behavior of the other player, we call it »L1«. Note that for a self-interested column player the strategy Center can also be due to L2 behavior. Similarly, the actions Top and Left are called »Nash« because they comprise all strategies that reflect higher levels than L1 for the row player as well as higher levels than L2 for the column player, including Nash play, under the assumption of self-interested preferences. As we are interested in learning as the amount of switching from the set of low-level to higher-level strategies, this rough classification is sufficient. Likewise, for a clear distinction between the considered preference types, we name the actions Bottom and Right »Utilitarian«, since they are consistent with all levels of reasoning higher than L1 (including Nash) under the assumption of utilitarian preferences. Therefore, when we describe some event as »an increase in Nash play«, we mean that we observe an increase in actions that are consistent with higher-level strategies, given self-interested preferences. Accordingly, a »decrease in L1« denotes a reduction of low-level strategic play, no matter which preferences are considered.

2.3 Treatments

To study the impact of information on choices and belief statements, we implemented four treatments, the details of which are given in Table 3. Our main interest is in the random-matching treatment, denoted by RM. In this treatment subjects had all relevant information about the game, i.e. the set of players, the set of strategies and the payoff function of each player. In addition, after each period they received feedback about the payoff earned in this period (and thereby about the action of the other player). In all treatments subjects did not receive any feedback about their payoffs from the belief elicitation task.⁶ In the second treatment, we only changed the matching

⁶Nevertheless, they could infer their payoff from this task after receiving feedback about the outcome of the game. The main reason for not showing the payoffs from the belief elicitation task was to change as few parameters as possible when going from RM, PI and FM to NF.

Treatment	Payoff	Feedback	Matching	Periods	Sessions	# of subjects
RM	own+opponent	own payoff	random	20	4	54
FM	own+opponent	own payoff	fixed	20	4	54
PI	own	own payoff	fixed	20	4	48
NF	own+opponent	none	fixed	20	4	50

Table 3: Treatments.

scheme to fixed matching, denoted by FM, in order to understand the role of the matching protocol for learning.

The two remaining treatments serve to control for the effect of information on the learning process. In treatment NF (no feedback), subjects received no feedback at all, but had common knowledge of the payoff structure of the game as in the baseline treatment. In treatment PI (partial information), subjects were only informed about their own payoff function, but not about the payoff function of their opponent. However, they received feedback after each period, such that they could infer the choice of their opponent.⁷

We conducted both treatments NF and PI with fixed matching. In treatment NF without any feedback about the behavior of the other player, the matching protocol does not matter for the game-theoretic prediction. We therefore compare the results from treatment NF to the baseline treatment with random matching. In treatment PI where we are interested in how players learn to play a game about which they only hold incomplete information, we employed fixed matching to keep the environment as simple as possible. Accordingly, we compare the results of PI to treatment FM. Note that repeated-game effects are in principle only possible in treatment FM, but not in RM, NF and PI. Without feedback in NF or without information about the payoffs of the other player in PI, strategies that punish a player for deviations from the equilibrium path are impossible.

2.4 Matching, beliefs and payments

At the beginning of a session, subjects were randomly assigned a player role (row player or column player), which they kept during the whole experiment. However, they made all their decisions

⁷We use the names *L1*, Nash and Utilitarian also in treatment PI even though a priori the subjects cannot reason about the other player's incentives and consequently cannot identify the Nash and the Utilitarian action. However, subjects can use their received feedback to construct a "subjective game". Kalai and Lehrer (1993) show that subjective games can converge to an ε -Nash equilibrium of the underlying game.

from the perspective of the row player, i.e. for column players we used a transformation of the matrix game in Table 1. Before choosing an action (choice task), we asked subjects to indicate their beliefs regarding the behavior of their opponent (belief task). In particular, we asked subjects to state the expected frequencies of play, i.e., they had to specify in how many out of 100 times they expected the column player to choose Left, Center and Right in the current period.⁸ After the belief task, subjects had to make their choice by selecting one of the three possible actions (mixing was not possible). We employed belief elicitation in all four treatments to analyze the impact of the matching scheme and information on beliefs and choices.

Subjects were paid for both tasks. For the choice task, we paid subjects according to the numbers in the payoff matrix, which were exchanged at the commonly known rate of 1 point = € 0.15. To reward the belief task, we used a quadratic scoring rule (QSR) which is incentive compatible given that subjects are risk-neutral money maximizers. The QSR we used is defined as follows. The payoff Π_{it}^{QSR} for player i in period t for a given action a_{jt}^k with $k \in \{L, C, R\}$ of player j in period t and belief vector $b_{it} = (b_{it}^L, b_{it}^C, b_{it}^R) \in \Delta^2$ such that $\Delta^2 = \{b_{it} \in \mathbb{R}^3 \mid \sum_{k \in \{L, C, R\}} b_{it}^k = 1\}$ is

$$\Pi_{it}^{QSR}(b_{it}, a_{jt}) = A - B * \left(\sum_{k \in \{L, C, R\}} (b_{it}^k - 1_{[a_{jt}^k]})^2 \right) \quad (1)$$

where $1_{[a_{jt}^k]}$ is an indicator function equal to 1 if a_{jt}^k is chosen in period t and 0 otherwise. While paying subjects for the choice and the belief task is necessary to ensure incentive compatibility, it allows subjects to engage in hedging. Subjects can for example coordinate on a cell of the payoff matrix that is not an equilibrium and become unwilling to move away from it in order to avoid losses in the belief task. To eliminate such behavior, we decided to determine the final payoffs as follows.⁹ First, at the end of the experiment we selected one period randomly and independently to determine the payoffs for each of the two tasks. Second, we used parameters $A = 1.5$ and $B = 0.75$ in the QSR. Thus, the maximum payoff from the belief task (€ 1.50) was relatively low compared to payoffs from choice task. For instance, the Nash equilibrium [Top, Left] would lead to payoffs of € 11.7 and € 10.2 for the two player roles.¹⁰

⁸For simplicity we restricted the expected frequencies of play to integers. Therefore, we count any belief statement assigning a weight of 34 percent to one action and 33 percent to each of the remaining actions as a uniform belief statement.

⁹Blanco et al (2008) propose and test a slightly different method to avoid hedging. Their hedging-proof method suggests paying randomly either the decision task or the belief elicitation task. They find no evidence for hedging.

¹⁰Note that subjects could guarantee themselves a payoff of € 1 by stating uniform beliefs. Although this would be an attractive choice for a risk-averse subject, we find no evidence of such behavior in our treatments. Only 7.2

The experiments were conducted in the computer lab at Technical University Berlin using the software tool kit *z-Tree*, developed by Fischbacher (2007). Subjects were recruited via a mailing list through which they could voluntarily register to participate in decision experiments (Greiner, 2004). Upon entering the lab, subjects received written instructions and were asked to read them carefully.¹¹ After everybody had finished reading the instructions, we distributed an understanding test that covered both the game and the QSR. Only after all subjects had answered the questions correctly, we proceeded with the experiment. In total 206 students (115 males and 91 females) from various disciplines participated in the four treatments. Sessions lasted about one hour. Subjects' average earnings were about € 12.80, including a show-up fee of € 3 for arriving at the laboratory on time.

3 Results

In the first part of the analysis, we examine the choices made by the experimental subjects. We begin this analysis with a focus on first period behavior and a comparison of these results to previous experiments. Afterwards we extend our analysis to all periods and focus on the development of behavior over time, considering the impact of the information available. In the second part of the data analysis, we make use of the elicited beliefs. After confirming that the stated beliefs outperform beliefs constructed with standard models of belief formation, we examine the frequency of best responses to the stated beliefs. Furthermore we check the accuracy of the stated beliefs in predicting the opponent's choice as well as the role of feedback and payoff information for the formation of beliefs.

Note that unlike in most other studies on asymmetric one-shot games (e.g. Costa-Gomes and Weizsäcker 2008), we do not pool the data over player roles. As we study only one specific game, we are able to consider the exact strategic situation of each player role. This differentiation would be lost by pooling the data. Thus, we run all statistical tests separately for row and column players. All results reported as significant in the paper are based on a 5%-level of significance.

percent of belief statements assign no less than 30 and no more than 35 percent to all three of the opponent's actions. (RM 5.2%, FM 5.8%, PI 5.9% and NF 12.1%)

¹¹For a sample of the instructions see the Appendix.

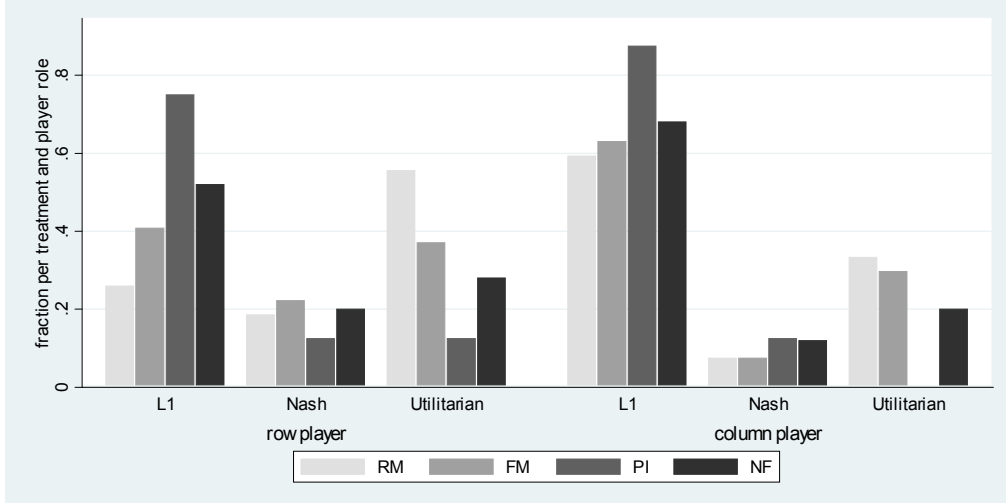


Figure 1: First period choices.

3.1 Choices

3.1.1 First-period choices

In this section, we look at behavior in the first period only. This is of some stand-alone interest, since many experiments on behavior in one-shot 3x3 normal-form games have used similar games, and we can compare our results to them. First-period play in our experiment differs from one-shot experiments because players know that they will play the game again. But according to the game-theoretic prediction, this should not affect play, with the exception of the fixed-matching treatment.

First-period behavior in the four treatments is presented in Figure 1. The figure shows the fraction of each action for all four treatments. In the first period, subjects in treatments RM, FM and NF are in a comparable situation, and we do not observe any differences in behavior, as can be taken from Figure 1. We cannot reject the hypothesis that the frequency of choices is the same in these three treatments using a χ^2 -Test.¹²

Excluding treatment PI where players face a different game and pooling the data over player roles, we observe 51% *L1* behavior in the first period in RM, FM and NF. This is in line with previous studies. For instance, Costa-Gomes et al. (2001) estimated a frequency of *L1* choices of about 45%, Rey-Biel (forthcoming) found 48% *L1* behavior in his constant-sum games, whereas Costa-Gomes and Weizsäcker (2008) observed slightly higher rates of about 60%.

¹²For both player roles we perform a pairwise comparison of RM with FM and NF, respectively. The test yields no p-value smaller than 0.1 ($\chi^2_{(2)}$).




Figure 2: Choices in all periods.

Now consider the decision situation in the first period of treatment PI. Subjects only know their own payoffs in the game and therefore cannot base their decisions on strategic considerations. Hence, it is no surprise to see 39 out of 48 subjects (81%) choosing the *L1* rule in period 1 in PI, which not only maximizes the minimum payoff, but also the expected payoff assuming that the opponent randomizes uniformly over all possible actions. Concerning the column player's choice of the dominated action Right (Utilitarian), it is remarkable that no column player in PI chooses the Utilitarian action in the first period. This indicates that the choice of dominated actions in the other treatments is due to the payoff structure of the other player and not to mistakes. The frequency of the three strategies in PI is significantly different from FM in the first period for both player roles ($\chi^2_{(2)}$, $p = 0.043$ for row players and $p = 0.014$ for column players). We summarize the findings on choices in the first period in the following result.

Result 1 (i) *First-period behavior in RM, FM and NF is statistically indistinguishable from each other and comparable to findings from one-shot experiments.* (ii) *Except for the row player in RM, L1 is the most frequently chosen strategy in the first period in all treatments and for both player roles.* (iii) *In treatment PI, the fraction of subjects choosing L1 in the first period is higher than in all other treatments.*

3.1.2 Choices over all periods

We now turn to the behavior in all 20 periods. First consider the proportion of the three actions averaged over all rounds, displayed in Figure 2. To compare the proportion of choices over all

	Row Player			Column Player		
	L1	L2+/Nash	Utilitarian	L1/L2	L3+/Nash	Utilitarian
<i>Const</i>	-0.98*** (0.20)	-0.73*** (0.20)	-0.09 (0.22)	0.41** (0.17)	-1.55*** (0.19)	-0.92*** (0.22)
<i>D_{FMPI}</i>	-0.27 (0.28)	0.14 (0.29)	0.07 (0.31)	-0.73*** (0.24)	0.35 (0.27)	0.46 (0.31)
<i>D_{PI}</i>	0.82*** (0.29)	0.31 (0.29)	-1.05*** (0.32)	0.67*** (0.25)	0.55** (0.27)	-1.21*** (0.33)
<i>D_{NF}</i>	0.57** (0.28)	-0.19 (0.29)	-0.35 (0.31)	-0.07 (0.25)	0.45* (0.27)	-0.21 (0.32)
<i>logL</i>	-976.4	-1045.1	-1042.0	-1205.2	-851.5	-844.4
$\chi^2_{(k-1)}$	12.8***	4.9	13.2***	12.0***	11.5***	13.9***
<i>N</i>	2060					

Panel-probit regression with random individual effects, standard errors in parentheses

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

Table 4: Choices on average (relative to RM and FM).

periods in the different treatments, we perform a regression for each strategy and player role combination. We regress the strategies on treatment dummies. This gives us a first indication of the influence of the different information conditions. To model the repeated decisions of the same subject in each treatment, we use a panel regression with random individual effects. Since subjects had to choose one out of three possible strategies, a probit model is employed where the dependent variable reflects the inclination to choose one strategy over the other two.

The results of the regression, summarized in Table 4, reveal the relative treatment effects on the proportion of choices. The difference between random and fixed matching is captured by the coefficient of D_{FMPI} .¹³ While the choices of row players are not affected by the matching protocol in a statistically significant way, column players choose *L1* less often with fixed than with random matching. The higher proportion of strategic choices with fixed matching can be explained by the simpler learning environment with a fixed partner. Note that repeated-game effects that could also account for differences between RM and FM would only affect the difference between the proportion

¹³The dummy D_{FMPI} is coded as 1 both for treatment FM and PI. With the separate dummy for PI, D_{PI} , we can thereby compare FM to PI.

	Row Player			Column Player		
	L1	L2+/Nash	Utilitarian	L1/L2	L3+/Nash	Utilitarian
$D_{RM} \cdot Period$	-0.04*** (0.01)	0.04*** (0.01)	-0.01 (0.01)	0.01 (0.01)	0.02* (0.01)	-0.04*** (0.01)
$D_{FM} \cdot Period$	-0.04*** (0.01)	0.03*** (0.01)	0.00 (0.01)	-0.02* (0.01)	0.04*** (0.01)	-0.01 (0.01)
$D_{PI} \cdot Period$	-0.04*** (0.01)	0.03*** (0.01)	0.00 (0.01)	-0.05*** (0.01)	0.05*** (0.01)	0.01 (0.02)
$D_{NF} \cdot Period$	-0.01 (0.01)	0.01 (0.01)	0.01 (0.01)	-0.01 (0.01)	0.02 (0.01)	-0.01 (0.01)
$log\mathcal{L}$	-958.7	-1030.6	-1041.4	-1192.0	-833.4	-837.8
$\chi^2_{(k-1)}$	77.0***	55.8***	25.6***	36.2***	110.9***	82.8***
N	2060					

Panel-probit regression with individual random effects, standard errors in parentheses estimated constants for each treatment have been omitted in the table

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

Table 5: Choices over time.

of Nash and Utilitarian play, not the proportion of $L1$ choices, on the equilibrium path.

Next we compare the baseline treatment RM with the no-feedback treatment NF with the help of the dummy D_{NF} . The lack of feedback in NF results in overall more $L1$ play than in RM for the row player, which can be ascribed to the no-feedback environment being less conducive to learning. The effect of information about the other player's payoff is captured by the coefficient D_{PI} (for the difference between FM and PI). In the partial-information treatment, there is significantly more $L1$ play and less Utilitarian play for both player roles as well as more Nash play of the column player. The lack of information about the other player's payoffs increases non-strategic choices, and the proportion of Utilitarian play becomes negligible as the Utilitarian outcome cannot be identified.

Next we turn to the development of behavior over time. To give a first impression of how subjects play the game in the different treatments, Figure 3 presents the evolution of choices for each treatment. The figure shows averages over three periods in a given treatment for row players in the left panel and for column players in the right panel. To investigate the potential learning paths, we use regressions with a time trend. The results of these regressions are presented in Table



Figure 3: Decision rules over time.

5. The findings are rather clear-cut: Non-strategic play ($L1$) decreases in all treatments and for both player roles for treatment NF and except for the column player in RM and FM. Secondly, Nash play increases in all treatments and for both player roles except for treatment NF and only marginally for the column player in RM. Thus, in the sense of Stahl and Wilson we observe a trend towards more strategic play (that is more Nash and less $L1$ play) in all treatments with feedback information. There is no indication that subjects learn to play Nash simply by introspection and thinking. In all treatments, the proportion of Utilitarian play hardly varies over time, the exception being the column player in RM.

The findings based on the various regressions can be summarized as follows.

Result 2 *(i) The level of non-strategic play is lowest in FM for both player roles. (ii) The proportion of non-strategic $L1$ play decreases in all treatments except NF. (iii) The proportion of Nash choices increases at least marginally in all treatments and for both player roles except in NF. (iv) The proportion of Utilitarian choices is almost constant over time for all treatments and player roles (except for the column player in RM).*

In PI, the overall lower proportion of strategic behavior compared to FM can be ascribed to the lack of information about the opponent's payoffs. However, the fact that players in PI can observe the choices of their opponent and react to these observations leads to a trend away from the $L1$ rule, just as in RM and FM. In treatment NF, behavior does not change over time. As the NF treatment is comparable to a repeated one-shot situation, this finding lends support to the frequently applied method of giving no feedback between different tasks in experiments in order to minimize learning effects.

Finally, treatment FM and RM are statistically indistinguishable for the row player. But we observe that the column player's behavior is affected by the matching protocol in that she chooses more non-strategic $L1$ play in RM than in FM. This difference can be ascribed to the fact that the column player's Utilitarian action is dominated and is thus chosen less often in the stranger design of RM than in FM with a partner design where column players "invest" in the cooperative outcome.

3.2 Belief formation

In this section, we focus on the relationship between the elicited beliefs and the subjects' own as well as their opponents' actions. In standard equilibrium analysis it is assumed that subjects form

beliefs about the behavior of the opponent and then best respond to these beliefs. The level- k model departs from this view by positing that subjects differ in their strategic sophistication when thinking about the behavior of other players, i.e., they differ in their beliefs (Stahl and Wilson, 1995). In particular, level-1 behavior implies that beliefs are naive in that uniform randomization by the opponent is assumed. Level-2 types hold the belief that others best respond to uniform randomization. Thus, we can use belief statements to measure the level of strategic sophistication and to track the development of strategic thinking over time.

At this point, we would like to address some caveats concerning elicited beliefs. First, subjects need not hold beliefs about the opponent's play at all. For example, they might choose some non-strategic decision rule in the first period and then condition play on received payoffs (as in reinforcement learning). Forcing them to state beliefs could alter the choices if these subjects move their decisions in the direction of belief-based play.¹⁴ However, our design is based on a comparison between treatments which all use belief elicitation. Unless the effects of belief elicitation interact with our treatment variables, our results are immune to such problems.

Second and more importantly, the assumption of best-responses to beliefs in decision theory can be understood as an "as if" assumption. With this interpretation, subjects do not necessarily have to best respond to their stated beliefs as these beliefs might be unrelated to the true underlying beliefs. In order to address this concern, we compare the stated beliefs to beliefs constructed from previous play of the opponent in the next subsection. The stated beliefs emerge as a better predictor of actual choices than the constructed beliefs, which lends support to the hypothesis that the elicited beliefs are the best approximations of the true underlying beliefs that are available.¹⁵

Third, even though we asked explicitly to state myopic beliefs, i.e. beliefs only for the current period, we cannot rule out that subjects follow repeated-game strategies and hold beliefs consistent with this. As the choices that are part of repeated-game strategies are not necessarily best responses to myopic beliefs, we expect best-response rates to be lower in FM than in RM if repeated-game strategies play a role.

¹⁴See Rutström and Wilcox (2006) for an argument along these lines.

¹⁵See Costa-Gomes and Weizsäcker (2008) for a thorough analysis of belief statements and their relationship to the true beliefs.

3.2.1 Stated beliefs vs. models of belief formation

We follow the approach used in Nyarko and Schotter (2002) and compare the explanatory power of elicited beliefs compared to standard belief learning models. The purpose of this comparison is to establish whether stated beliefs are a good measure of strategic uncertainty or whether stated beliefs are inferior to beliefs derived indirectly from the opponents' choices.

Standard belief learning models assume that players update their beliefs based on the opponent's history of play and then best-respond to these beliefs. The two most prominent models based on this assumption are the fictitious-play and the Cournot best-response model. While in the Cournot model subjects best respond to the opponent's play in the very last period, players in a pure fictitious-play model best respond to beliefs based on all previous actions of the opponent. The γ -weighted fictitious-play model introduced by Cheung and Friedman (1997) contains Cournot best response and fictitious-play as special cases. In this model subject i 's belief $b_{i,t+1}^k$ that subject j will choose action $a_{jt}^k, k \in \{L, C, R\}$ in period $t + 1$ is defined as:

$$b_{i,t+1}^k = \frac{1[a_{jt}^k] + \sum_{u=1}^{t-1} \gamma^u 1[a_{j,t-u}^k]}{1 + \sum_{u=1}^{t-1} \gamma^u}. \quad (2)$$

The parameter γ is the weight the player gives to the past actions of his opponent. It is obvious from (2) that $\gamma = 0$ leads to the Cournot best-response model and $\gamma = 1$ yields fictitious-play, respectively. We incorporate this model into a standard logistic choice model to allow subjects to best respond to their beliefs with noise. Subject i chooses action k with probability

$$\Pr(a_{it}^k | b_{it}) = \frac{\exp(\lambda \pi[a_{it}^k, b_{it}])}{\sum_{l \in \{L, C, R\}} \exp(\lambda \pi[a_{it}^l, b_{it}])}, \quad (3)$$

where $\pi[a_{it}^k, b_{it}]$ is the expected payoff of player i when she chooses an action k given her beliefs b_{it} over the action set of her opponent. The parameter λ determines the impact of this expected payoff on her own choice probability and can be interpreted as a rationality parameter. A player with $\lambda = 0$ chooses all actions with equal probability disregarding the expected payoff of her choice. On the other hand if $\lambda \rightarrow \infty$ the player is fully rational, i.e. she always best responds to her beliefs.

With respect to the specification of individual preferences, the preceding analysis on choices has demonstrated that information about the other player's payoffs leads to a significant increase in Utilitarian play. This supports the hypothesis that the payoffs of others may matter for an individual's utility. To avoid misspecification, we incorporate this finding in the following analysis by allowing for other-regarding preferences

$$u(m, y) = m(a_{it}^k, a_{jt}^k) + \theta y(a_{it}^k, a_{jt}^k), \quad (4)$$

which is identical to the basic preference model of Cox et al. (2007) under the assumption of risk neutrality. In the model, $m(a_{it}^k, a_{jt}^k)$ denotes the player's own payoff given the actions of both players whereas $y(a_{it}^k, a_{jt}^k)$ denotes the corresponding payoff of the other player. Thus θ is the willingness to exchange own for other's payoff which, in the case of risk neutrality, is equal to the marginal utility of an additional unit of the other player's payoff ($WTP = 1/MRS = (\partial u/\partial y)/(\partial u/\partial m) = \theta$). For $\theta = 0$, expression (4) turns into self-interested preferences as used e.g. by Nyarko and Schotter (2002).

We now turn to the estimation and probabilistic comparison of the choice model (3) based on the γ -weighted fictitious-play model (2) on the one hand, and based on the stated beliefs on the other hand. The model assumes that subjects process information about their own and the other's payoffs as well as about the history of the other's play. Because of the latter, we have to exclude treatment NF in the estimation. Besides the FM data, we use the data from treatment RM, since the process described in (2) can also be interpreted as the formation of beliefs over the average play of the population rather than over individual choices. Furthermore we can run the regression using the data from treatment PI with θ being restricted to 0. The estimation results for each treatment and player role are presented in Table 6.¹⁶

Treatm	Role	ML-estimation of model (3) using							Model selection	
		Fictitious beliefs (2)				Stated beliefs			Vuong (1989)	
		λ	θ	γ	$\log\mathcal{L}$	λ	θ	$\log\mathcal{L}$	Z	p-value
RM	Row	0.017*	0.337	0.575***	-575.5	0.090***	0.060	-488.9	-6.5	0.000
	Col	0.072***	-0.105	0.905***	-462.0	0.066***	-0.067	-425.0	-2.6	0.010
FM	Row	0.052***	0.132	0.682***	-480.7	0.104***	0.212***	-390.0	-4.9	0.000
	Col	0.031***	0.684	0.670***	-478.6	0.076***	0.615***	-375.9	-6.3	0.000
PI	Row	0.044***	-	0.649***	-487.7	0.065***	-	-451.2	-3.7	0.000
	Col	0.057***	-	0.622***	-413.6	0.107***	-	-308.0	-5.8	0.000

p-values are two-sided. Clarke's (2007) test gives the same results and yields very similar p-values.

Table 6: Comparison of fictitious and stated beliefs.

As a first result we observe that the belief models play a significant role in explaining the behavior of our subjects. This is especially for the model using the stated beliefs, since here the

¹⁶For the γ -weighted fictitious-play model we estimated γ and λ simultaneously. The ML-estimations and tests have been conducted with Stata and Matlab.

hypothesis that the rationality parameter λ is equal to zero is rejected for all treatments and player roles. Using tests for the selection between non-nested models introduced by Vuong (1989) and Clarke (2003), the hypothesis of equal explanatory power of the models can be rejected at all usual significance levels for all treatments and player roles, whereas the stated belief model is always closer to the real data generating process than the belief-learning model.¹⁷

To summarize, we extend the finding of Nyarko and Schotter (2002) from a matching-pennies game to our normal-form game with a unique subgame perfect Nash equilibrium in pure strategies. Furthermore, we allow for other-regarding preferences and find evidence for them. The estimations show that stated beliefs are better at explaining observed choices than beliefs that are implied by the standard models of belief formation. In the following, we therefore use the stated beliefs when analyzing the impact of experience and information on the consistency and accuracy of beliefs.

3.2.2 Consistency of actions and stated beliefs

Both in standard Nash equilibrium and in the level-k model it is assumed that subjects best respond to their beliefs. Using the elicited beliefs, we can investigate the consistency of actions and stated beliefs, i.e. whether subjects best respond to their stated beliefs. This helps us to evaluate the relative descriptive validity of assuming best-response behavior in the four different treatments. However, in treatment FM the possibility of repeated-game strategies implies that subjects do not necessarily choose a best response to their myopic belief. For example, if column players in FM expect the row player to choose the Utilitarian action and respond by choosing it as well for the sake of keeping up cooperation in later periods, this choice does not represent a best response. Thus, best-response rates will be lower in FM than in RM if repeated-game effects play a role.

In Figure 4 the proportion of players best responding to their stated beliefs is displayed for each player role and treatment separately. The figure shows the average proportion of best responses over three periods. In all treatments, the average best-response rates are rather low, ranging mainly from 45% to 75%. In order to compare our results to other studies, it is useful to

¹⁷Vuong's test statistic is based on the overall likelihood ratio of two rival models and is asymptotically normally distributed under the null. Clarke's test statistic consists of the number of single likelihood ratios being greater than 1 which is binomially distributed under the null with parameters $\theta = 0.5$ and the number of observations in each subset of the data. Vuong's test is outperformed by Clarke's test when the distribution of the single log-likelihood ratios is highly peaked. Both tests were calculated using corrections for the dimension of the models as proposed by Schwarz (1978) and Clarke (2007) respectively.

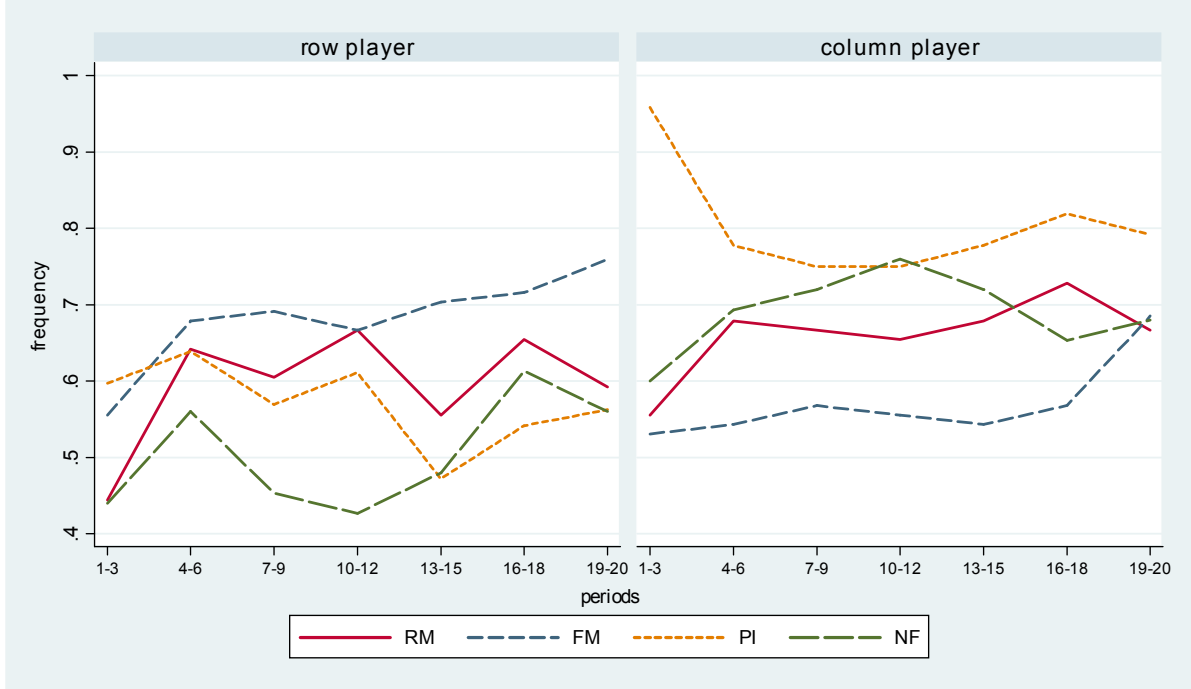


Figure 4: Best-response rates over time.

look at the aggregated best-response behavior of all subjects. Averaging over all treatments and player roles, subjects best-respond to their stated beliefs in 63% of the cases (in RM in 63% of the cases). This is in line with best-response rates found in similar studies. In simple games like 2x2 games (Nyarko and Schotter, 2002) or constant-sum games (Rey-Biel, forthcoming), consistency rates are about 70%, whereas the rates range from 49% to 63% in more complicated games like ours or the games used in Costa-Gomes and Weizsäcker (2008) and Ehrblatt et al.(2008), respectively.

For statistical evidence on differences between the treatments and the development of best-response rates over time, we run random-effects panel regressions. As the dependent variable is either 0 (no best response) or 1 (best response), we use a probit model. Besides the constant, the independent variables are dummies for FM/PI, PI and NF. In addition, we test for a linear time trend in each treatment. The regression results are summarized in Table 7 and 8.

Differences in the level of best-response rates between treatments are not very large, as displayed in Table 7. They only exception is the column player in PI who best responds more often than in FM simply because he rarely chooses the Utilitarian action that he cannot identify as such.

This finding of no strong differences in best-response rates between treatments is in line with the theory according to which best-response behavior is independent of the information players have. However, when considering whether subjects learn to best respond in the baseline treatment RM

Best Response Rates		
	Row Player	Column Player
<i>Const</i>	0.30** (0.15)	0.56*** (0.19)
<i>D_{FMPI}</i>	0.25 (0.21)	-0.39 (0.26)
<i>D_{PI}</i>	-0.36* (0.22)	0.85*** (0.27)
<i>D_{NF}</i>	-0.26 (0.21)	0.12 (0.27)
<i>logL</i>	-1237.8	-1058.1
$\chi^2_{(k-1)}$	6.2	10.2**
<i>N</i>	2060	

Panel-probit regression with random individual effects.

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

Table 7: Best-response rates.

in the course of the experiment, the significant and positive coefficient of Period in Table 8 reveals that this is the case for the column player and also marginally for the row player. Similarly, in FM best-response rates increase significantly for the row player, and marginally significantly for the column player. The two other treatments, PI and NF, do not display significant increases in best-response rates.

This raises two questions. First, why do best-response rates increase at all? Second, why do best-response rates increase in treatments RM and FM, but not in PI and NF? Internal consistency requires best responding to one's beliefs, independent of the information conditions and a player's experience with a game. The results from treatments RM and FM suggest that learning to play a game seems to encompass learning to be internally consistent. In treatment NF, however, subjects might be doubtful about the accuracy of their beliefs, lacking information about the other player's behavior. This might induce them to put less weight on their beliefs when choosing an action. But this reasoning fails to explain the similar result in treatment PI where there is also no discernible increase in best-response behavior. In PI, players have to learn about the structure of the game over

Best Response Rates		
	Row Player	Column Player
$D_{RM} \cdot Period$	0.02*	0.02**
	(0.01)	(0.01)
$D_{FM} \cdot Period$	0.03***	0.02*
	(0.01)	(0.01)
$D_{PI} \cdot Period$	-0.01	-0.02*
	(0.01)	(0.01)
$D_{NF} \cdot Period$	0.02*	0.01
	(0.01)	(0.01)
$\log \mathcal{L}$	-1229.2	-1052.6
$\chi^2_{(k-1)}$	32.3***	50.1***
N	2060	

Panel-probit regression with random individual effects, estimated constants for each treatment have been omitted in the table

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

Table 8: Best-response rates over time.

time. Thus, the complexity of learning the structure of the game and learning to best respond to one's beliefs at the same time may be too high. Second, in treatment PI many subjects start with uniform beliefs and best respond to them. As the belief set of $L1$ is large and $L1$ is an attractive strategy initially, there is a high rate of consistency at the outset. This effect is absent in RM, FM and NF.

The focus of the preceding analysis was on myopic beliefs. In the repeated-game setting of treatment FM, folk theorem results are possible. If subjects aim at a cooperative outcome, column players might choose their dominated action (Utilitarian) when expecting Utilitarian play of row players. This explains why we observe lower best-response rates and more Utilitarian choices for column players in FM compared to NF and PI. But we observe a substantial proportion of Utilitarian play also in RM in both player roles. Moreover, the regressions reveal no significant differences between FM and RM neither for the overall proportion of Utilitarian play (see Table 4), nor for the average best response rates (see Table 7). This suggests that it is mainly the subjects'

preference for maximizing the overall payoff that leads to the high level of Utilitarian actions of the row player in FM, RM and NF, not repeated-game effects.

The insignificant difference of best-response rates in FM and RM could be due to a higher number of failures to best respond to undominated actions in RM, which would push best-response rates down in the direction of FM. But this is not the case. When considering only the best-response behavior to Nash and $L1$, we find best response rates of about 92% in FM and 88% in RM. We can further support this finding of equal best-response rates in FM and RM by a Kolmogorov-Smirnov test which compares the number of best responses to Nash and $L1$ of each subject. The test yields a p-value of $p > 0.88$.¹⁸ For these reasons we consider the evidence for repeated-game strategies as weak.

Result 3 (i) *The overall level of best responses does not differ significantly between treatments RM and FM as well as RM and NF. Only column players in PI best-respond significantly more often than in FM.* (ii) *While the proportion of best responses increases over time in RM and FM, there is no significant time trend in NF and PI.*

Interpreting the stated beliefs as proxies for the true underlying beliefs, we can conclude that actors learn to best respond more often to their beliefs in games with feedback information and information about the game structure with some experience of the situation, compared to situations with less information and experience. Thus, actors become more sophisticated over time in that the consistency of their actions and beliefs increases. This is a novel observation, and we are not aware of any model of rational or boundedly rational choice which can account for this finding.

3.2.3 Accuracy of stated beliefs

We will now focus on whether the elicited beliefs are accurate in predicting the behavior of the opponents. As the accuracy of beliefs is a measure of strategic sophistication, it differs under the Nash equilibrium prediction and under the level-k model. In the Nash equilibrium of the stage game, subjects hold accurate beliefs about their opponent's choice. In the level-k model, however, this is not necessarily the case as subjects' beliefs can be at odds with their opponents' behavior. The experimental data from the different treatments allow us to identify the factors enabling subjects

¹⁸We use each column player as an independent observation and compare the empirical distribution of the number of best responses to Nash and $L1$ between FM and RM.




Figure 5: Accuracy of stated beliefs.

to predict their opponent’s play and to state accurate beliefs. In order to measure how well stated beliefs predict the opponent’s play, we use the earnings from the quadratic scoring rule (QSR).

Figure 5 shows the earnings (averaged over three periods) from the QSR for all treatments and for both player roles.¹⁹ The average payoff across treatments and player roles is about € 1.²⁰ This corresponds to the payoff for a subject who states uniform beliefs, which is indicated by the vertical line in Figure 5. The second benchmark to which we can compare the earnings is € 0.50, representing the expected payoff from randomizing uniformly over degenerate beliefs.

Although subjects earned hardly more than € 1, their beliefs were more accurate than if they simply tried to predict the choice of their opponent with a probability of one (Wilcoxon signed-rank test, all p-values < 0.01). For row players in all four treatments, we cannot reject the hypothesis of equal means at a 5% level of significance for all treatments (Wilcoxon signed-rank test, p-values > 0.085). The same holds for column players in FM and PI, but column players in NF and RM earned on average significantly less than € 1 (Wilcoxon signed-rank test, for NF and

¹⁹In principle, the accuracy of predicting other’s behavior should not depend on the player role. Indeed, we only find a weakly significant difference between player roles in RM (Mann-Whitney test, $p = 0.051$). In all other treatments the same test yields p-values higher than 0.49.

²⁰The average payoff across player roles is € 0.92 in RM, € 1.07 in FM, € 1.02 in PI and € 0.89 in NF.

Belief Accuracy		
	Row Player	Column Player
<i>Const</i>	0.96*** (0.04)	0.88*** (0.04)
<i>D_{FMPI}</i>	0.09 (0.06)	0.20*** (0.06)
<i>D_{PI}</i>	-0.02 (0.06)	-0.07 (0.06)
<i>D_{NF}</i>	-0.05 (0.06)	-0.02 (0.06)
<i>logL</i>	-1403.8	-1386.4
$\chi^2_{(k-1)}$	7.4*	18.7***
<i>N</i>	2060	

Panel-probit regression with random individual effects.

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

Table 9: Accuracy of stated beliefs.

RM p -values < 0.01).

Our main interest again lies in the development over time. If players become more strategic in the course of the experiment and reason more about the incentives of the opponent, the accuracy of beliefs should increase. Notice that this interpretation encompasses cases where play converges to a combination of choices, and players therefore hold correct beliefs. Figure 5 displays such improvements over time in predicting the opponent’s play in all treatments except for NF. For the statistical analysis, we ran a random-effects panel regression where the dependent variable is the payoff from the belief elicitation task. The results are displayed in Table 9. In addition to the constant, the regression includes treatment dummies for the controls FM/PI, PI and NF as independent variables. The only significant difference concerns the column player who exhibits a lower accuracy of beliefs in RM than in FM. Again, this can be explained by the higher predictability of a fixed partner.

In addition, we performed tests of the time trends in all treatments reported in Table 10. Here, we observe a clear pattern. The column players improve their predictions in all treatments

Belief Accuracy		
	Row Player	Column Player
$D_{RM} \cdot Period$	0.00 (0.00)	0.01** (0.00)
$D_{FM} \cdot Period$	0.01 (0.00)	0.01*** (0.00)
$D_{PI} \cdot Period$	0.01* (0.00)	0.01*** (0.00)
$D_{NF} \cdot Period$	-0.01* (0.00)	0.00 (0.00)
$log\mathcal{L}$	-1398.9	-1366.8
$\chi^2_{(k-1)}$	1642.2***	1621.6***
N	2060	

Panel-probit regression with random individual effects, estimated constants for each treatment have been omitted in the table.

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

Table 10: Accuracy of stated beliefs over time.

except in NF. For the row players there is no significant time trend. The findings can be summarized as follows:

Result 4 (i) Overall, there is no significant difference between the accuracy of beliefs in the four treatments when comparing RM with FM and with NF, as well as FM with PI. Only column players in FM submit more accurate beliefs than in RM. (ii) In treatments RM, FM and PI, behavior is characterized by a similar learning path in that the column player's beliefs are more accurate in later periods, while there is no time trend for the row player.

The results indicate that feedback about past behavior of one's opponent(s) is more important for learning to predict choices than information about the full game. In addition, playing with the same opponent facilitates accurate predictions of choices. Thus, feedback information and fixed matching allow players to form accurate beliefs which are an important ingredient of Nash equilibrium play.

4 Summary and Conclusions

We performed an experiment to study the development of strategic reasoning over a limited number of periods. To classify the strategies of the 3x3 normal-form game employed in our study, we used the level-k model of Stahl and Wilson (1995) and allowed both for selfish and other-regarding preferences. This classification of choices allowed us to track strategic play over time. In order to understand the determinants of strategic play, we varied the information available to the players and elicited their beliefs about opponents' play.

We find that feedback information and information about the payoffs of the opponent have an impact on choices. When either type of information is lacking, this leads to an increase in non-strategic ($L1$) play. The absence of information about the opponent's payoffs additionally leads to a decrease in Utilitarian play. However, not revealing the opponent's payoff function has almost no impact on the learning path. In all treatments except for NF, subjects exhibit less non-strategic and more Nash play over time. In contrast, in the no-feedback treatment there is no increase in strategic play in the course of the experiment. This fact clearly highlights the importance of feedback and the limits of deductive reasoning of the subjects.

Regarding the analysis of beliefs, we first evaluate whether stated beliefs or beliefs constructed with belief-learning models are a better proxy for the underlying true beliefs of the subjects. We find that the stated beliefs are more consistent with actual choices than beliefs constructed with belief models such as weighted fictitious play or Cournot best response. Given this result, we study the best-response rates to the stated beliefs. In treatments RM and FM, actions are consistent with stated beliefs more frequently in later periods. Incomplete information about the opponent's payoff function in PI or no feedback in treatment NF inhibit this trend towards more best responses in later periods.

The accuracy of the subjects' beliefs with respect to the opponent's choices is increasing over time in all treatments except in NF. Remarkably, removing the information about the opponent's payoff function does not significantly decrease the overall accuracy of beliefs nor its development over time. However, without feedback information players are not able to improve their predictions of the other player's behavior in the course of the experiment.

Summing up, the higher proportion of non-strategic choices in terms of the level-k model and the lower accuracy of beliefs in early compared to later rounds in treatments PI, FM and also in RM show that subjects learn to play the game in environments with feedback, but independent of

the information on the payoffs of the other player and the matching protocol. This validates the use of inductive learning models. Moreover, we find an increase in the internal consistency of choices and beliefs in the course of the experiment, but only in the treatments with full information about the game and feedback. This finding is by now very little understood and in our view deserves more thorough empirical scrutiny.

References

- [1] Blanco, M., Engelmann, D., Koch, A. and Normann, H. (2008). Belief Elicitation in Experiments: Is there a Hedging Problem?, Working Paper.
- [2] Camerer, C., Ho, T., and Chong, J.-K. (2004). A Cognitive Hierarchy Model of Games, *Quarterly Journal of Economics*, 119, 861-898.
- [3] Clarke, K. A. (2003). Nonparametric Model Discrimination in International Relations, *Journal of Conflict Resolution*, 47, 72-93.
- [4] Clarke, K. A. (2007). A Simple Distribution-Free Test for Nonnested Model Selection, *Political Analysis*, 15, 347-363.
- [5] Costa-Gomes, M.A., Crawford, V. and Broseta, B. (2001). Cognition and Behavior in Normal-form Games: An Experimental Study, *Econometrica*, 69, 1193-1235.
- [6] Costa-Gomes, M.A. and Weizsäcker, G. (2008). Stated Beliefs and Play in Normal-form Games, *Review of Economic Studies*, 75, 729-762.
- [7] Cox, J. C., Friedman, D. and Gjerstad, S. (2007). A tractable model of reciprocity and fairness, *Games and Economic Behavior*, 59, 17 - 45.
- [8] Crawford, V. and Iriberri, N. (2007a). Level-k Auctions: Can a Non-Equilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private Value Auctions, *Econometrica*, 75, 1721-1770.
- [9] Crawford, V. and Iriberri, N. (2007b). Fatal Attraction: Salience, Naivete, and Sophistication in Experimental Hide-and-Seek Games, *American Economic Review*, 97, 1731-1750.
- [10] Ehrblatt, W. Z., Hyndman, K., Özbay, E.Y. and Schotter, A. (2008). Convergence: An Experimental Study of Teaching and Learning in Repeated Games, Working Paper.
- [11] Fischbacher, Urs (2007). z-Tree: Zurich toolbox for ready-made economic experiments, *Experimental Economics*, 10, 171-178.
- [12] Gneezy, U. (2005). Step-Level Reasoning and Bidding in Auctions, *Management Science*, 51, 1633-1642.

- [13] Greiner, B. (2004). An Online Recruitment System for Economic Experiments, in: Kremer, K., Macho, V. (eds.): *Forschung und wissenschaftliches Rechnen 2003*, GWDG Bericht 63, Göttingen: Ges. für Wiss. Datenverarbeitung, 79-93.
- [14] Ivanov, A. (2006). Strategic Play and Risk Aversion in One-Shot Normal-Form Games: An Experimental Study, Working Paper.
- [15] Kalai, E. and Lehrer, E. (1993). Subjective Equilibrium in Repeated Games, *Econometrica*, 61, 1231-1240.
- [16] Nyarko, Y. and Schotter, A. (2002). An Experimental Study of Belief Learning Using Elicited Beliefs, *Econometrica*, 70, 971-1005.
- [17] Oechssler, J. and Schipper, B. (2003). Can You Guess the Game You Are Playing?, *Games and Economic Behavior*, 43, 137-152.
- [18] Rey-Biel, P. (forthcoming). Equilibrium Play and Best Response to (Stated) Beliefs in Constant Sum Games, *Games and Economic Behavior*.
- [19] Rutström, E. and Wilcox, N. (2006). Stated Beliefs Versus Empirical Beliefs: A Methodological Inquiry and Experimental Test. Mimeo.
- [20] Schwarz, G. (1978). Estimating the Dimension of a Model, *The Annals of Statistics*, 6, 461-464.
- [21] Stahl, D.O. and Wilson, P.W. (1995). On Players' Models of Other Players: Theory and Experimental Evidence. *Games and Economic Behavior*, 10, 218-254.
- [22] Terracol, A. and Vaksman, J. (2009). Dumbing Down Rational Players: Learning and Teaching in an Experimental Game. *Journal of Economic Behavior and Organization*, 70, 54-71.
- [23] Vuong, Q. H. (1989). Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses, *Econometrica*, 57, 307-333.
- [24] Weber, R. (2003). 'Learning' With No Feedback in a Competitive Guessing Game, *Games and Economic Behavior*, 44, 134-144.
- [25] Weber, R. and Rick, S. (2008). Meaningful learning and transfer of learning in games played repeatedly without feedback, Working Paper.

Appendix

Instructions (for FM)

The experiment you are about to participate in is part of a project financed by the German Research Foundation (DFG). Its aim is to analyze economic decision-making behavior. You can earn a considerable amount of money in this experiment, dependent on your decisions and the decisions of the other participants. Consequently, it is extremely important that you read these instructions very carefully.

Please note: these instructions are for your eyes only, and it is not permitted to hand on any information whatsoever to other participants. Similarly, you are not allowed to speak to the other participants throughout the whole experiment. Should you have a question, please raise your hand and we will come to you and answer your question individually. Please do not ask your question(s) aloud. If you break these rules, we will unfortunately be compelled to discontinue the experiment.

General information The experiment is made up of several periods where decisions must be made and questions answered. You can win points with your decisions. These points represent your earnings and will be converted into euros at the end of the game and paid out in cash. The exact procedure of the experiment, the various decisions and the method of payment are clearly explained in the next section.

The decision-making situation At the beginning of the experiment, you will be assigned by draw to another participant, randomly and anonymously. This allocation is maintained throughout the whole of the remaining experiment. The participant who has been assigned to you will be called “the other one” from now on.

In each period, you and the other one will be confronted with the same decision-making situation. Each time, you must choose between the three alternatives: “top”, “middle”, and “bottom”.

Each of these three alternatives has been given three possible payoffs (as points). The other one must also decide between three alternatives (“left”, “center” or “right”), and each of these alternatives has also three possible payoffs, as above. You will see the following input screen on the computer:

round

1 out of 20

remaining time [sec]: 30

	Decision of the other one: Left	Decision of the other one: Center	Decision of the other one: Right
Your Decision: Top	⁶⁸ 78	²³ 72	²⁰ 12
Your Decision: Middle	⁵² 67	⁶³ 59	⁴⁹ 78
Your Decision: Bottom	¹¹ 21	⁸⁹ 62	⁷⁸ 89

Your Decision: Top Middle Bottom

Next

Your three alternatives, “top”, “middle”, and “bottom”, are listed in the first column of the table. Next to your alternatives, you can see three boxes, each with two numbers. The subscript (lower) number is always your possible payoff. On the input screen illustrated above, the alternative “top” has been allocated the payoff of 78, 72 and 12, the alternative “middle” the payoff of 67, 59 and 78, and the alternative “bottom” the payoff of 21, 62 and 89. This means that should you decide on “top”, for example, then your payoff is 78, 72 or 12 points. The payoff you actually receive depends on whether the other one selects “left”, “center” or “right”. Thus your payoff depends on your own decision as well as that of the other one. The superscript (raised) number in any box is always the possible payoff of the other one. For example, if the other one decides on “left”, then his/her possible payoff points are 68, 52 and 11. This means, for example, that if you decide on “middle” and the other one decides on “right”, your payoff is 78 points. The payoff for the other one is 49 points in this case.

The possible payoff points on the input screen above are therefore as follows:

You choose “top”; the other one chooses “left”:	
Your payoff is:	78 points
The payoff for the other one is:	68 points
You choose “top”; the other one chooses “center”	
Your payoff is:	72 points
The payoff for the other one is:	23 points
You choose “top”; the other one chooses “right”:	
Your payoff is:	12 points
The payoff for the other one is:	20 points
You choose “middle”; the other one chooses “left”:	
Your payoff is:	67 points
The payoff for the other one is:	52 points
You choose “middle”; the other one chooses “center”:	
Your payoff is:	59 points
The payoff for the other one is:	63 points
You choose “middle”; the other one chooses “right”:	
Your payoff is:	78 points
The payoff for the other one is:	49 points
You choose “bottom”; the other one chooses “left”:	
Your payoff is:	21 points
The payoff for the other one is:	11 points
You choose “bottom”; the other one chooses “center”:	
Your payoff is:	62 points
The payoff for the other one is:	89 points
You choose “bottom”; the other one chooses “right”:	
Your payoff is:	89 points
The payoff for the other one is:	78 points

Please note that the possible payoff points for you and the other one remain the same in every period.

The other one always has exactly the same input screen in front of him/her as you do. After you and the other one have chosen between the three alternatives, you will be informed of your

payoff in this period. This is the only information you will be given during the experiment in each period. The next period begins after that.

Statement of expectations

a) How can you state your expectations? Before each decision-making situation, you will be asked how you estimate the decision-making behavior of the other one. This means that at the beginning of each period we will require you to predict how the other one will decide in this period. You will have to answer the following question:

In how many out of 100 cases do you expect the other one to decide on “left”, “center” or “right”?

Of course, the other one makes his decision only once in each period. You could also consider the question as asking you to state the likelihood that each of the three alternatives is chosen by the other one. You will see the following input screen on the computer:

round
1 out of 20
remaining time [sec]: 30

	Decision of the other one: Left	Decision of the other one: Center	Decision of the other one: Right
Your Decision: Top	68 78	23 72	20 12
Your Decision: Middle	52 67	63 59	49 78
Your Decision: Bottom	11 21	89 62	78 89

In how many out of 100 cases do you expect the other one to decide on "left", "center" or "right"?

Left

Center

Right

Your three alternatives, “top”, “middle” and “bottom”, are listed in the table above, as well as the corresponding possible payoff. Below that, there is the question with the three boxes.

Let us assume that you are sure that the other one will choose “right”, and definitely not “center” or “left”. Then you would respond to our question by entering the number 100 in the box for “right” and the number 0 in the boxes for “center” and “left”. Alternatively, we could assume that you think the other one will probably choose “center”, but there is still a small chance that s/he will choose “right”, and an even smaller chance that s/he will choose “left”. Then, for example, you might respond to our question by entering the number 70 for “center”, 20 for “right” and 10 for “left”.

If you think it is even more unlikely that s/he will choose “center”, then you could enter, for example, 60 for “center”, 24 for “right” and 16 for “left”. Or it is possible that you think it is equally likely that the other one will choose “left”, “center” and “right”. Then you should enter, for example, the numbers 33, 33, 34 in the boxes.

Please note that the three numbers may not be decimal, and that they must always add up to 100.

N.B.: The numbers used in the examples have been chosen arbitrarily. They give you no indication how you and the other one decide.

b) How is the payoff for your stated expectations calculated? Your payoff is calculated after you have guessed how frequently the other one chooses his/her three alternatives. Your payoff depends on the difference between your estimate of the frequency of the decision and the actual decision made. Your payoff is higher when you have guessed that the other one often makes the “true” decision (which s/he really made), and it is lower when you have guessed that the other one will make this decision infrequently. Similarly, your payoff is higher when you have correctly predicted that the other one will not make a particular decision and then s/he in fact does not make the decision.

The exact calculation of the payoff is as follows: We calculate a number for each of the three alternatives. This number reflects how appropriate your estimate of the decision frequency of the corresponding alternative was. We take these three numbers to calculate your payoff.

First, we consider how well you predicted the alternatives which were actually chosen. Let us assume that the other one chose “left”. We then compare your estimate of how often the other one would choose “left” out of 100 cases with the number 100, and calculate the difference between

the two. This difference is then multiplied by itself and the resulting number multiplied by the factor 0.0005. Thus, if you expected the other one to choose “left” in many out of 100 cases, then this number will be smaller (since the difference between your estimate and 100 is small) than if you expected that s/he would choose “left” in few out of 100 cases.

Then we consider how well you predicted that the other two alternatives would not be chosen. Let us assume again, for example, that the other one chose “left”, which at the same time means that “center” and “right” were not chosen. Then we take your estimate for the alternative “center” and multiply this by itself. The resulting number is again multiplied by the factor 0.0005. We apply this procedure again to your estimate for the alternative “right”. We then take the three numbers thus calculated and deduct them from the number 10. This determines the number of points you receive for your statement of expectations.

As an illustration of how your payoff might appear, let us consider three examples. Let us assume that the other one chose “left” and that your estimate for “left” was 100 and correspondingly 0 for the other two alternatives. This means that you have stated an estimate that is exactly right. Consequently, you earn the following points:

$$10 - 0.0005 * (100 - 100)^2 - 0.0005 * 0^2 - 0.0005 * 0^2 = 10$$

Let us assume again that the other one chose “left”. Your estimate for “left” was 60, for “center” 20 and for “right” 20, which means that your stated estimate predicted that the other one would choose “left” more frequently than “center” and “right”. Consequently, you earn the following points:

$$10 - 0.0005 * (100 - 60)^2 - 0.0005 * 20^2 - 0.0005 * 20^2 = 8.8$$

If we still assume that the other one chose “left”, but your estimate for “left” was 0, for “center” also 0 and for “right” 100, this means that your stated estimate was exactly wrong. Consequently, you earn the following points:

$$10 - 0.0005 * (100 - 0)^2 - 0.0005 * 0^2 - 0.0005 * 100^2 = 0$$

N.B.: The numbers used in the examples have been chosen arbitrarily. They give no indication how you and the other one decide.

These examples should make it clear that you will always receive a payoff of at least 0 points, and at most 10 points for your stated expectations. And the closer your estimations, the more money you earn. (You may be asking yourself why we have chosen such a payoff ruling as described above. The reason being that with such a payoff ruling, you can expect the highest payment when you state numbers that are closest to your own estimate.)

Procedure and payment The experiment consists of 20 periods altogether. In each period, you have to first state your estimate of the behavior of the other one, and then make your own decision.

At the end of the experiment, a period each for the decision-making situation and for the statement of expectations will be chosen randomly in order to determine your earnings in the experiment. The choice of both periods will be made randomly by the experiment leader throwing a dice. The chosen periods will then be entered onto the input screen by the experiment leader. At the end of the experiment, you will see an overview of your earnings from the decision-making situation and your earnings from the statement of expectation, as well as the total amount. The payoff that you have attained in the corresponding period chosen will be converted at a rate of

1 point = 15 cents

and will be paid out in cash.

Do you have any questions?

Control questions Now you have to answer 7 questions. In this way we are checking whether you have understood the decisions you have to make during the experiment. Should you have any further questions, please raise your hand and one of the experiment leaders will come to you. The experiment will not start until all participants have answered the control questions correctly.

The decision-making situation:

round
1 out of 20
remaining time [sec]: 30

	Decision of the other one: Left	Decision of the other one: Center	Decision of the other one: Right
Your Decision: Top	68 78	23 72	20 12
Your Decision: Middle	52 67	63 59	49 78
Your Decision: Bottom	11 21	89 62	78 89

Your Decision:
 Top
 Middle
 Bottom

Next

1. If you choose “bottom” and the other one chooses “center”, how many points do you earn?

2. If you choose “middle” and the other one chooses “left”, how many points does the other one earn?

3. If we assume your payoff amounts to 12, which decision did the other one make?

4. If you choose “bottom” and the other one chooses “left”, how much do you earn and how much does the other one earn?

The other one: _____ You: _____

5. Consider the following two cases:

You expect the other one to choose “left” in 80 out of 100 cases. The other one actually does choose “left”. You expect the other one to choose “left” in 20 out of 100 cases. The other one actually chooses “right”. In both cases we assume that you expect the other one to choose “center” in 0 out of 100 cases.

Is your payoff for the statement of expectation in the first case:

higher the same lower (Please underline your answer!)

than in the second case?

6. Imagine that Participant 1 states the following expectation: The other one chooses “left” in 50 out of 100 cases, “center” in 20 out of 100 cases, and “right” in 30 out of 100 cases. Participant 2 expects the following: the other one chooses “left” in 60 out of 100 cases, “center” in 20 out of 100 cases, and “right” in 20 out of 100 cases. We will assume that the other one chose “left” by Participant 1 as well as by Participant 2. Who will receive the highest payoff?

Participant _____

7. If you consider all three alternatives to be equally possible, which numbers should you then enter?

left: _____ center: _____ right: _____

Thank you for participating in the experiment!

Bücher des Schwerpunkts Märkte und Politik
Books of the Research Area Markets and Politics

- Kai A. Konrad, Tim Lohse (Eds.)
Einnahmen- und Steuerpolitik in Europa: Herausforderungen und Chancen
2009, Peter Lang Verlag
- Kai A. Konrad
Strategy and Dynamics in Contests
2009, Oxford University Press
- Roger D. Congleton, Arye L. Hillman, Kai A. Konrad (Eds.)
40 Years of Research on Rent Seeking
2008, Springer
- Kai A. Konrad, Beate Jochimsen (Eds.)
Föderalismuskommission II: Neuordnung von Autonomie und Verantwortung
2008, Peter Lang Verlag
- Mark Gradstein, Kai A. Konrad (Eds.)
Institutions and Norms in Economic Development
2007, MIT Press
- Johannes Münster
Mobbers, Robbers, and Warriors
2007, Shaker Verlag
- Kai A. Konrad, Beate Jochimsen (Eds.)
Der Föderalstaat nach dem Berlin-Urteil
2007, Peter Lang Verlag
- Kai A. Konrad, Beate Jochimsen (Eds.)
Finanzkrise im Bundesstaat
2006, Peter Lang Verlag
- Robert Nuscheler
On Competition and Regulation in Health Care Systems
2005, Peter Lang Verlag
- Pablo Beramendi
Decentralization and Income Inequality
2003, Madrid: Juan March Institute
- Thomas R. Cusack
A National Challenge at the Local Level: Citizens, Elites and Institutions in Reunified Germany
2003, Ashgate
- Sebastian Kessing
Essays on Employment Protection
2003, Freie Universität Berlin
<http://www.diss.fu-berlin.de/2003/202>
- Daniel Krähmer
On Learning and Information in Markets and Organizations
2003, Shaker Verlag
- Tomaso Duso
The Political Economy of the Regulatory Process: An Empirical Approach
Humboldt-University Dissertation, 2002, Berlin,
<http://edoc.hu-berlin.de/dissertationen/duso-tomaso-2002-07-17/PDF/Duso.pdf>
- Bob Hancké
Large Firms and Institutional Change. Industrial Renewal and Economic Restructuring in France
2002, Oxford University Press
- Andreas Stephan
Essays on the Contribution of Public Infrastructure to Private: Production and its Political Economy
2002, dissertation.de
- Peter A. Hall, David Soskice (Eds.)
Varieties of Capitalism
2001, Oxford University Press
- Hans Mewis
Essays on Herd Behavior and Strategic Delegation
2001, Shaker Verlag
- Andreas Moerke
Organisationslernen über Netzwerke – Die personellen Verflechtungen von Führungsgremien japanischer Aktiengesellschaften
2001, Deutscher Universitäts-Verlag
- Silke Neubauer
Multimarket Contact and Organizational Design
2001, Deutscher Universitäts-Verlag
- Lars-Hendrik Röller, Christian Wey (Eds.)
Die Soziale Marktwirtschaft in der neuen Weltwirtschaft, WZB Jahrbuch 2001
2001, edition sigma
- Michael Tröge
Competition in Credit Markets: A Theoretic Analysis
2001, Deutscher Universitäts-Verlag
- Torben Iversen, Jonas Pontusson, David Soskice (Eds.)
Unions, Employers, and Central Banks
2000, Cambridge University Press
- Tobias Miarka
Financial Intermediation and Deregulation: A Critical Analysis of Japanese Bank-Firm-Relationships
2000, Physica-Verlag
- Rita Zobel
Beschäftigungsveränderungen und organisationales Lernen in japanischen Industriegesellschaften
2000, Humboldt-Universität zu Berlin
<http://dochost.rz.hu-berlin.de/dissertationen/zobel-rita-2000-06-19>
- Jos Jansen
Essays on Incentives in Regulation and Innovation
2000, Tilburg University

DISCUSSION PAPERS 2008

Dan Kovenock Brian Roberson	Inefficient Redistribution and Inefficient Redistributive Politics	SP II 2008 – 01
Dan Kovenock Brian Roberson	Coalitional Colonel Blotto Games with Application to the Economics of Alliances	SP II 2008 – 02
Vito Tanzi	The Future of Fiscal Federalism	SP II 2008 – 03
Kai A. Konrad Kjell Erik Lommerud	Love and Taxes – and Matching Institutions	SP II 2008 – 04
Benny Geys Jan Vermeir	Party Cues and Yardstick Voting	SP II 2008 – 05
Benny Geys Jan Vermeir	The Political Cost of Taxation: New Evidence from German Popularity Ratings	SP II 2008 – 06
Kai A. Konrad Dan Kovenock	The Alliance Formation Puzzle and Capacity Constraints	SP II 2008 – 07
Johannes Münster	Repeated Contests with Asymmetric Information	SP II 2008 – 08
Kai A. Konrad Dan Kovenock	Competition for FDI with Vintage Investment and Agglomeration Advantages	SP II 2008 – 09
Kai A. Konrad	Non-binding Minimum Taxes May Foster Tax Competition	SP II 2008 – 10
Florian Morath	Strategic Information Acquisition and the Mitigation of Global Warming	SP II 2008 – 11
Joseph Clougherty Anming Zhang	Domestic Rivalry and Export Performance: Theory and Evidence from International Airline Markets	SP II 2008 – 12
Jonathan Beck	Diderot’s Rule	SP II 2008 – 13
Susanne Prantl	The Role of Policies Supporting New Firms: An Evaluation for Germany after Reunification	SP II 2008 – 14
Jo Seldeslachts Tomaso Duso Enrico Pennings	On the Stability of Research Joint Ventures: Implications for Collusion	SP II 2008 – 15
Dan Kovenock Brian Roberson	Is the 50-State Strategy Optimal?	SP II 2008 – 16
Joseph Clougherty Tomaso Duso	The Impact of Horizontal Mergers on Rivals: Gains to Being Left Outside a Merger	SP II 2008 – 17
Benny Geys Wim Moesen	Exploring Sources of Local Government Technical Inefficiency: Evidence from Flemish Municipalities	SP II 2008 – 18
Benny Geys Friedrich Heinemann Alexander Kalb	Local Governments in the Wake of Demographic Change: Evidence from German Municipalities	SP II 2008 – 19
Johannes Münster	Group Contest Success Functions	SP II 2008 – 20
Benny Geys Wim Moesen	Measuring Local Government Technical (In)efficiency: An Application and Comparison of FDH, DEA and Econometric Approaches	SP II 2008 – 21

DISCUSSION PAPERS 2009

Aron Kiss	Coalition Politics and Accountability	SP II 2009 – 01
Benny Geys Friedrich Heinemann Alexander Kalb	Voter Involvement, Fiscal Autonomy and Public Sector Efficiency: Evidence from German Municipalities	SP II 2009 – 02
Salmai Qari Kai A. Konrad Benny Geys	Patriotism, Taxation and International Mobility	SP II 2009 – 03
Kai A. Konrad Salmai Qari	The Last Refuge of a Scoundrel? Patriotism and Tax Compliance	SP II 2009 – 04
Sven Chojnacki Nils Metternich Johannes Münster	Mercenaries in Civil Wars, 1950-2000	SP II 2009 – 05
Oliver Gürtler Johannes Münster	Sabotage in Dynamic Tournaments	SP II 2009 – 06
Dan Kovenock Brian Roberson	Non-Partisan ‘Get-Out-the-Vote’ Efforts and Policy Outcomes	SP II 2009 – 07
Subhasish M. Chowdhury Dan Kovenock Roman M. Sheremeta	An Experimental Investigation of Colonel Blotto Games	SP II 2009 – 08
Michael R. Baye Dan Kovenock Casper G. de Vries	Contests with Rank-Order Spillovers	SP II 2009 – 09
Florian Morath Johannes Münster	Information Acquisition in Conflicts	SP II 2009 – 10
Benny Geys	Wars, Presidents and Popularity: The Political Cost(s) of War Re-examined	SP II 2009 – 11
Paolo Buccirossi Lorenzo Ciari Tomaso Duso Giancarlo Spagnolo Cristiana Vitale	Competition policy and productivity growth: An empirical assessment	SP II 2009 – 12
Pedro P. Barros Joseph Clougherty Jo Seldeslachts	How to Measure the Deterrence Effects of Merger Policy: Frequency or Composition?	SP II 2009 – 13
Paolo Buccirossi Lorenzo Ciari Tomaso Duso Giancarlo Spagnolo Cristiana Vitale	Deterrence in Competition Law	SP II 2009 – 14
Paolo Buccirossi Lorenzo Ciari Tomaso Duso Giancarlo Spagnolo Cristiana Vitale	Measuring the deterrence properties of competition policy: the Competition Policy Indexes	SP II 2009 – 15

DISCUSSION PAPERS 2010

Dorothea Kübler	Experimental Practices in Economics: Performativity and the Creation of Phenomena	SP II 2010 – 01
Dietmar Fehr Dorothea Kübler David Danz	Information and Beliefs in a Repeated Normal-form Game	SP II 2010 – 02

Bei Ihren Bestellungen von WZB-Papers schicken Sie bitte unbedingt einen an Sie adressierten Aufkleber mit sowie je paper eine Briefmarke im Wert von 0,51 Euro oder einen "Coupon Reponse International" (für Besteller aus dem Ausland)

Please send a self addressed label and postage stamps in the amount of 0.51 Euro or a "Coupon-Reponse International" (if you are ordering from outside Germany) for each WZB-paper requested

Bestellschein

Order Form

Absender / Return Address:

Wissenschaftszentrum Berlin
für Sozialforschung
Presse- und Informationsreferat
Reichpietschufer 50

D-10785 Berlin-Tiergarten

**Hiermit bestelle ich folgende(s)
Discussion paper(s):**

**Please send me the following
Discussion paper(s):**

Bestell-Nr. / Order no.	Autor/in, Kurztitel /Author(s) / Title(s) in brief