

Ensthaler, Ludwig; Huck, Steffen; Leutgeb, Johannes

Article — Accepted Manuscript (Postprint)

Games played through agents in the laboratory — a test of Prat & Rustichini's model

Games and Economic Behavior

Provided in Cooperation with:

WZB Berlin Social Science Center

Suggested Citation: Ensthaler, Ludwig; Huck, Steffen; Leutgeb, Johannes (2020) : Games played through agents in the laboratory — a test of Prat & Rustichini's model, Games and Economic Behavior, ISSN 1090-2473, Elsevier, Amsterdam, Vol. 119, pp. 30-55, <https://doi.org/10.1016/j.geb.2019.09.013>

This Version is available at:

<https://hdl.handle.net/10419/218846>

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.



<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Ludwig Ensthaler
Steffen Huck
Johannes Leutgeb

Games played through agents in the laboratory – A test of Prat & Rustichini's model

Ludwig Ensthaler, Rocket Internet SE, Berlin
Steffen Huck, WZB and University College London (steffen.huck@wzb.eu)
Johannes Leutgeb, WZB (johannes.leutgeb@wzb.eu)

Corresponding author at: WZB, Reichpietschufer 50, 10785 Berlin, Germany.
E-mail address: steffen.huck@wzb.eu (S.Huck).

September 2019

Abstract

From the regulation of sports to lawmaking in parliament, in many situations one group of people ("agents") make decisions that affect the payoffs of others ("principals") who may offer action-contingent transfers in order to sway the agents' decisions. Prat and Rustichini (2003) characterize pure-strategy equilibria of such Games Played Through Agents. Specifically, they predict the equilibrium outcome in pure strategies to be efficient. We test the theory in a series of experimental treatments with human principals and computerized agents. The theory predicts remarkably well which actions and outcomes are implemented but subjects' transfer offers deviate systematically from equilibrium. We show how quantal response equilibrium accounts for the deviations and test its predictions out of sample. Our results show that quantal response equilibrium is particularly well suited for explaining behavior in such games.

Keywords: games played through agents, experiment, quantal response equilibrium.

JEL classification: D44, C91, D72, D83

1 Introduction

Consider an individual, A, whose payoff depends on how some agents play a game. This individual has no immediate choice to make in this game, yet its outcome might be of considerable importance to her. Imagine a pharma company whose business opportunities depend on a new piece of legislation to be voted on in parliament. Or a motor racing team whose competitive edge might depend on new rules to be introduced by the governing body of the sport. Or, perhaps, a child who has no direct say in her parents' relocation decision or choice of holiday destination but passionately cares about both. In all these cases A will mull over the consequences of the game played by the agents and her only option to influence the outcome is through some form of "bribery", that is, through providing incentives to the agents, who play the game, to take particular choices that may lead to one of her more desired outcomes.

Of course, at the same time, there might be another individual, B, who also cares about the outcome of this game and her incentives might differ from A's. B will then compete with A for influence over the agents. The situation that arises is essentially a bundle of auctions where A and B vie for influence by promising payments to the agents for making particular choices. A might offer a payment to some member of parliament for voting against the new legislation, while B might offer a payment to the same parliamentarian for voting for the new law. Of course, the parliamentarian might also have intrinsic preferences over the different outcomes that can ensue. Her total utility will then be a function of both, her intrinsic utility and the payments she can ensure for herself by voting for one of the actions that A and B have incentivized.

The full game that arises from this structure has the following timing. First, A and B make binding payment promises to all the agents. Then the agents play the game. This interesting and important class of games has been introduced by

Prat and Rustichini (2003) and termed *Games Played Through Agents* or GPTAs. We will stick to this terminology and will refer below to the players who play the actual game as the *agents* while we will refer to A and B as the *principals*. GPTAs are solved through backward induction. Under the assumption of additively separable utility, Prat and Rustichini (2003) study the equilibrium solutions for this class of games and find, somewhat surprisingly, that subgame perfect equilibria are characterized by strong efficiency. Specifically, equilibria yield efficient outcomes that maximize the sum of all players' payoffs, principals' and agents' alike.

Despite their wide applicability GPTAs have, to the best of our knowledge, not yet been subjected to an empirical investigation. Our study provides a first attempt to explore the empirical validity of Prat and Rustichini's (2003) theory by means of a laboratory experiment. As the equilibrium logic is far from being trivial we focus on some of the most basic cases, all based on 2×2 games where the agents who play the game have no intrinsic preferences over the four possible outcomes (perhaps such as some politicians who mainly value the income they obtain from lobbyists). Moreover, in order to avoid complications stemming from social preferences that could muddle the relation between agents and principals, we computerize the agents who simply implement the action that maximizes their income.

We study three types of GPTAs, where the payoff matrices, from the viewpoint of the principals, *resemble* the following games¹: a prisoners' dilemma (PD), a coordination game (COORD) and a simple dominance solvable game (DOM).² The PD and COORD games have symmetric payoff matrices while DOM is asymmetric. For the PD we implement two different payoff structures to test one of the theory's key

¹It is important to keep in mind that the base game that *the agents* play is always the same – with zeros in all cells.

²We also conducted another asymmetric treatment with a payoff matrix resembling a battle of the sexes, which is reported in Appendix A.

comparative statics which, incidentally, also includes predictions for the coordination game. While we find the equilibrium logic for the symmetric games already quite demanding, asymmetric games are even harder to solve due to the inherent conflict between principals.

Analyzing the data from our experiment we find that the theory does remarkably well. In each of our games the predicted actions are implemented most of the time and the comparative statics for equilibrium offers hold empirically across games and are highly significant. We did not necessarily expect this to be the case and were surprised by the theory's predictive power. Of course, the question arises how good is the fit really? What is the comparison? The point prediction of all probability mass concentrated on one outcome is, of course, falsified. So how can we say that the theory is doing a really good job organizing the data? On the one hand, we appeal to standard notions of good fit in experimental economics where hitting the target 80 % of the time tends to be a success. On the other hand, we examine the systematic deviations in offers from equilibrium predictions in more detail. We find a strong asymmetry in deviations: principals' offers are close to equilibrium for actions that matter most to them while offers for actions that are typically not implemented are too low. This is intuitive and can be shown to influence our subjects' reasoning from the very beginning of the game. In order to quantify these deviations in a systemic manner, we estimate a structural logit quantal response equilibrium (QRE) model. Our precision estimates are within the usual range and, for the symmetric games, the fit between QRE predictions and data is almost scarily good. To test the *predictive* power of our QRE framework out of sample we implement a further treatment, DOM. In the symmetric treatments QRE predicts an asymmetry *within* a principal's offers but in DOM, QRE predicts an asymmetry *between* two principals' offers. Our data exhibits the predicted pattern supporting the QRE approach. Our

results suggest that subjects violate the agent indifference condition in Prat and Rustichini's (2003) theory. This is a simple equilibrium condition that demands that agents are indifferent between both actions as otherwise bids ("bribes") could be reduced. Integrating QRE into the original theory allows us to explain the deviations and also successfully predict behavior out of sample.

In all, our paper shows that despite its non-trivial equilibrium logic, Prat and Rustichini's (2003) theory does succeed in the laboratory. This should encourage, both, further empirical studies of GPTAs (in the laboratory but perhaps also in the field) as well as more applications of the theory which so far have been rare.³

The remainder of the paper is organized as follows. Section 2 summarizes the key insights of Prat and Rustichini (2003). In Section 3 we introduce the symmetric games, discuss the experimental results and show how quantal response equilibrium organizes the data. In Section 4 we move on to the asymmetric games to test QRE and the theory further. Finally, Section 5 discusses the results and concludes.

2 Games Played Through Agents: A Very Short Summary

A Game Played Through Agents or GPTA, as defined in Prat and Rustichini (2003), is played by a set of agents who have to choose actions and a set of principals who offer each agent a schedule of monetary transfers contingent on chosen actions. An agent chooses her action to maximize the sum of transfers she receives from the

³The applications that we are aware of have dealt with lobbying in a federal (Bordignon et al., 2008) or international setting (Aidt and Hwang, 2008), where lobbyists take the role of principals and states or nations the role of agents. Fredriksson and Millimet (2007) apply Prat and Rustichini's (2003) theory to a game of pollution taxation and lobbying and argue that macro-level patterns in gasoline prices are consistent with the theory.

principals plus any intrinsic cost or benefit of her choice. A principal chooses her transfer schedule to maximize her utility from the agents' actions minus the sum of the transfers she makes to the agents. A GPTA is modeled as a two-stage game where, first, the principals simultaneously choose their transfer schedules and, second, all agents simultaneously choose their actions. Prat and Rustichini (2003) derive necessary and sufficient conditions for the existence of a pure-strategy equilibrium in these games with arbitrary numbers of agents, principals, and actions.

In this paper we only study games with two agents and two principals. Furthermore, each agent only has two actions to choose from and the agents derive no intrinsic utility from an outcome, that is, the base game that they play in the absence of interfering principals would be the game shown in table 1. In the presence of the two principals an agent's payoff from an action simply equals the sum of the promised transfers for taking that action. Let the game's notation be as follows. An agent $n \in N = \{R, C\}$ chooses an action $s_n \in S_n$, where $S_R = \{U, D\}$ and $S_C = \{L, R\}$. We will refer to the agents as row agent and column agent. The combination of agents' actions translates into an outcome $s \in S = \prod_{n \in N} S_n = \{UL, UR, DL, DR\}$. The principals $m \in M = \{A, B\}$ receive payoffs depending on the outcome. We will refer to the principals as principal A and principal B. Let π_s^m denote the gross payoff to principal m for outcome s . The principals' payoffs can then be represented in a matrix like table 1. Principals engage in a bidding game and offer transfers contingent on actions to agents. Let $t_n^m(s_n)$ denote the transfer principal m offers to agent n for choosing action s_n . If the agents implement outcome s principal m receives the net payoff $\pi_s^m - \sum_{n \in N} t_n^m(s_n)$. Agents do not derive intrinsic utility from the different outcomes of the game and, thus, simply choose the action that implies the highest total transfer. Framing the problem differently, the principals engage in two simultaneous sealed-bid first-price auctions where their valuation for winning each auction

depends on whether they win or lose the other auction.⁴

		L	R		L	R	
	U	0, 0	0, 0		U	π_{UL}^A, π_{UL}^B	π_{UR}^A, π_{UR}^B
	D	0, 0	0, 0		D	π_{DL}^A, π_{DL}^B	π_{DR}^A, π_{DR}^B
Agents' Base Game				Principals' Gross Payoffs Depending on Agents' Actions			

Table 1: Gross Payoff Matrices for Agents and Principals

Prat and Rustichini (2003) derive a particularly simple equilibrium characterization for this class of games. They show that the following three conditions characterize a pure-strategy equilibrium (\hat{t}, \hat{s}) : Agent Indifference (AI), Incentive Compatibility (IC), and Cost Minimization (CM).

$$\sum_{m \in M} \hat{t}_n^m(\hat{s}_n) = \sum_{m \in M} \hat{t}_n^m(s'_n) \quad \forall s'_n \in S_n, \forall n \in N \quad (\mathbf{AI})$$

Agent Indifference says that agents are indifferent between the equilibrium action \hat{s}_n and the alternative s'_n . Transfer offers are the same for both actions as otherwise some principal could decrease her transfer offer and implement the same outcome. Notice that while agents are indifferent, in equilibrium, they will implement the efficient outcome as otherwise the equilibrium logic would unravel.

$$\pi_{s'}^m - \pi_{\hat{s}}^m \leq \sum_{n \in N} \hat{t}_n^{j \neq m}(\hat{s}_n) - \sum_{n \in N} \hat{t}_n^{j \neq m}(s'_n) \quad \forall m \in M, \forall s' \in S \quad (\mathbf{IC})$$

Incentive Compatibility states that the gains from implementing a different outcome are outweighed by the costs of incentivizing the agents to do so. On the left-hand

⁴For the case of two principals and two agents who implement equilibria in a 2×2 subgame, that is, the case we study in the laboratory, GPTAs coincide with Bikhchandani (1999)'s auctions of heterogeneous objects. For larger GPTAs this equivalence does not hold as in a GPTA principals' valuations do not only depend on the set of auctions that they win but also on the identity of the winner of those auctions that they do not win themselves.

side of the inequality is the gain in gross payoffs for principal m from moving from \hat{s} to s' . On the right-hand side is the difference in the other principal's total offers for the actions that implement \hat{s} and s' respectively, or in other words, the additional costs that have to be borne in order to incentivize the agents to switch to another outcome. Remember that the matrix shown in table 1 only shows the principals' payoffs that result from the 2×2 game *played by the agents*—it is *not* the game played by the principals. In a GPTA, a principal must also compare payoffs moving along the diagonal (or off-diagonal) as she has to compare losing both auctions to winning both auctions.

$$\hat{t}_n^m(\hat{s}_n) > 0 \Rightarrow \hat{t}_n^m(s'_n) = 0 \quad \forall m \in M, \forall n \in N \quad (\mathbf{CM})$$

Finally, Cost Minimization requires that if a principal offers a positive transfer for the equilibrium action \hat{s}_n , she will not do so for the alternative s'_n . Else a principal could decrease her offers for both actions and still implement the same outcome. From these conditions it is easy to see that the equilibrium outcome must be efficient in the strong sense that in equilibrium the *sum* of payoffs to principals and agents will be maximal. Simply, sum (IC) over all principals and by (AI) the right-hand side equals zero. Therefore, the sum of gains from implementing a different outcome than \hat{s} cannot be positive, hence, the equilibrium outcome maximizes the sum of payoffs. In the following, we will refer to this simply as the efficient outcome.⁵

$$\sum_{m \in M} \pi_{s'}^m - \sum_{m \in M} \pi_{\hat{s}}^m \leq \sum_{m \in M} \sum_{n \in N} \hat{t}_n^m(\hat{s}_n) - \sum_{m \in M} \sum_{n \in N} \hat{t}_n^m(s'_n) = 0$$

⁵The analysis in Prat and Rustichini (2003) extends to more general classes of games. Also, they derive results for the case where the transfers of the principals are not contingent on the actions of the individual agents but on the outcome of the game which is a result of the actions of all agents. We do not discuss these results here as they are beyond the scope of our experimental analysis.

L R	L R	L R
U 4,4 0,6	U 4,4 0,5	U 3,3 0,0
D 6,0 1,1	D 5,0 1,1	D 0,0 1,1
PD High	PD Low	COORD

Table 2: Symmetric Gross Principals’ Payoff Matrices Depending on Agents’ Actions

3 The Experiment: Symmetric Settings

3.1 Equilibrium Derivation

We implement three symmetric GPTAs. The first game, PD High, uses a prisoners’ dilemma payoff matrix with a high temptation payoff. The second game, PD Low, does the same but with a lower temptation payoff. Finally, the third game, COORD, uses a coordination game payoff matrix, in which the principals’ incentives are perfectly aligned. Table 2 summarizes the payoff matrices for principals in the three treatments. In order to avoid negative payoffs we also introduce budget constraints equal to the maximum gross payoff that is attainable. Therefore, in PD High each principal’s budget is 6, in PD Low it is 5, and in COORD it is 3. In Prat and Rustichini’s (2003) model principals can offer transfers for all the agents’ actions but will, in equilibrium, only incentivize one of each. For the purposes of the experiment, we simplify the setup by restricting principals to offer transfers only for one of each agent’s actions, namely the one they prefer. In the symmetric treatments we restrict principal A to offer transfers for L and D, and conversely principal B may offer transfers for R and U. So $t_R^A(D), t_C^A(L), t_R^B(U), t_C^B(R) \geq 0$ and $t_R^A(U), t_C^A(R), t_R^B(D), t_C^B(L) = 0$. Consequently, **(CM)** is satisfied by design in all treatments.

PD High If a pure-strategy equilibrium exists it must have UL as an outcome, as this is the unique efficient outcome. By **(AI)**, the equilibrium offers are of the form $t_R^A(D) = t_R^B(U)$ and $t_C^A(L) = t_C^B(R)$ and hence both principals make both agents in-

different between choosing either of their actions. Restating the incentive constraints **IC** we know that

$$4 + t_C^A(L) \geq \max\{t_R^A(D) + t_C^A(L), 6, 1 + t_R^A(D)\},$$

and

$$4 + t_R^B(U) \leq \max\{t_R^B(U) + t_C^B(R), 6, 1 + t_C^B(R)\}$$

which simplifies to

$$2 \leq t_R^A(D) = t_R^B(U) \leq 4 \text{ and } 2 \leq t_C^A(L) = t_C^B(R) \leq 4.$$

As both agents have a budget of 6, we arrive at the following set of pure-strategy equilibria with UL as the outcome for PD High.

$$(t_R^A(D) = t_R^B(U), t_C^A(L) = t_C^B(R)) \in \{(a, b) \mid a, b \in [2, 4] \wedge a + b \leq 6\}$$

Intuitively, think of principal A who considers an outcome UL that promises her a gross payoff of 4. First, if the row agent were to choose D and thus implement DL principal A would receive a gross payoff of 6 for a net gain of 2. Thus, as long as principal B offers at least 2, principal A has no incentive to offer more. Instead she can match principal B's offer as, in equilibrium, the agent will choose U and she will not have to pay her offer. Second, consider that principal A were to decrease her offer to the column agent and would consequently lose the auction for this agent. Her gross payoff would then fall by 4. Therefore, as long as principal B offers at most 4 to the column agent, principal A will offer the exact same amount and will have to pay the offer for the column agent. Third, in order to implement DR she would have

to lose the auction for L while winning the auction for D. From the first two cases we know that from losing the former she would gain at most 4 while winning the latter would cost at least 2. The net gains in offers would amount to at most 2, while the net loss in gross payoffs would be 3. Thus, the conditions from the first two cases imply the conditions for the third. Fourth, by symmetry, the same considerations apply to principal B. Therefore equilibrium offers must be between 2 and 4. Since the agents are computerized, we expect principals to coordinate on the equilibrium with the lowest transfers to agents. We shall refer to

$$t_R^A(U) = t_R^B(D) = t_C^A(L) = t_C^B(R) = 2$$

as the (unique) lowest-transfer equilibrium.

PD Low By the same reasoning as in PD High we find that the set of pure-strategy equilibria with UL as the outcome for PD Low is

$$(t_R^A(U) = t_R^B(D), t_C^A(L) = t_C^B(R)) \in \{(a, b) \mid a, b \in [1, 4] \wedge a + b \leq 5\}$$

Again, there exists a unique lowest-transfer equilibrium where

$$t_R^A(U) = t_R^B(D) = t_C^A(L) = t_C^B(R) = 1$$

COORD As in PD High and PD Low, UL is the only efficient outcome and hence must be the outcome in any pure-strategy equilibrium. The set of pure-strategy equilibria with UL as the outcome is

$$(t_R^A(U) = t_R^B(D), t_C^A(L) = t_C^B(R)) \in \{(a, b) \mid a, b \leq 0 \wedge a + b \leq 2\}$$

The unique lowest-transfer equilibrium in this game is

$$t_R^A(U) = t_R^B(D) = t_C^A(L) = t_C^B(R) = 0$$

3.2 Experimental Implementation and Hypotheses

In order to reduce complexity and focus on the most interesting part of the model we computerize the two agents in all treatments. This also eliminates potential effects from social preferences between agents and principals which is realistic in many real-life settings where principals and agents might come from very different walks of life. As all games that we implement are 2×2 games, one agent is always called the row agent and the other agent the column agent. Moreover, for the symmetric games discussed above we chose to have identical instructions for all participants. That is, every principal received instructions as if they were principal A matched up against a principal B. When matching two principals, the computer then transformed one of the principal's offers into that of an equivalent principal B. Hence, offers for U and L (D and R) are equivalent and we can pool them into U/L (D/R).

Subjects are assigned the role of one of the two principals. Each subject receives an initial budget equal to the maximum payoff in the payoff matrix. The subjects can use this budget to incentivize the agents. In order to simplify matters further principals can only incentivize one of the two actions for each agent. (Remember that, in equilibrium, it would never make sense to offer an agent payments for both possible actions.) For example, principal A can only incentivize the row agent to choose D and the column agent to choose L. Agents will then choose the action for which they are offered the higher amount. In case both principals offer an agent the same amount, we need, of course, a tie-breaking rule in the experiment. As discussed above,

in equilibrium agents are necessarily indifferent (see Agent Indifference condition in section 2) and implement the efficient outcome. Hence, theory provides a tie-breaker for equilibrium offers. The equilibrium logic forces them to implement the efficient outcome as, if they did not, the equilibrium would unravel. This leaves the issue of which tie-breaking rule to implement when offers tie but are off equilibrium. To keep things simple in the experiment we decided to use the same tie-breaking rule for all cases as separate rules might confuse subjects or draw their attention to the equilibrium offers. The tie-breaking rule simply states that if the column agent is offered the same amount of money by both principals she will choose L and that if the row agent is offered the same amount of money by both agents, she will choose U. In the experiment just 2 % of tied offers are off equilibrium.

Let us now state some hypotheses for play in the three symmetric games that we expect to hold in the experiment. First, we would expect subjects to post offers such that the agents implement U and L.

Hypothesis 1. *U and L are the modal implemented actions in PD High, PD Low and COORD.*

Next notice that the offers for $(U/L, D/R)$ in the unique lowest-transfer equilibrium decrease from PD High to PD Low to COORD, from $(2, 2)$ to $(1, 1)$ to $(0, 0)$, respectively. Hence, from a comparative static perspective we would expect to see offers decrease from PD High to PD Low to COORD.

Hypothesis 2. *Median offers decrease from PD High to PD Low to COORD.*

Finally, all equilibria predict symmetric offers.

Hypothesis 3. *Principals choose symmetric offers in each of the three games.*

Procedures, player matching and all other details of the implementation are the same in treatments PD High, PD Low, and COORD and the instructions are fully analogous.⁶ There are 30 rounds of the same game in every session and matching between rounds is random. After each round subjects receive feedback about the offers of their opponent, the resulting outcome and their net earnings. The first two rounds are practice rounds while the remaining 28 rounds are payoff relevant. At the end of the experiment, two rounds out of those are chosen with equal probability and the earnings in those rounds are paid out to subjects. For each symmetric treatment we run two experimental sessions with 24 subjects each. Subjects are randomly assigned to matching groups of eight subjects.⁷ Hence, there are 48 subjects in six disjoint matching groups for every treatment. Subjects are matched within these matching groups for the duration of the experiment. In every round, a subject is randomly assigned to one of the other seven subjects in her matching group, throughout the experiment with the same uniform probability. Before the start of the first round subjects have to take an understanding test (see appendix D which also contains the instructions). Subjects are only allowed to continue with the experiment after answering all questions in the test correctly. Altogether, we run five treatments, with the three symmetric games described above as well as two asymmetric setups described further below. The sessions were run between 2014 and 2016 at the WZB-TU laboratory in Berlin with subjects who are undergraduate or master's students at one of the major research universities in town. In total 264 subjects participated in the experiment. The experiment was implemented in z-Tree (Fischbacher, 2007) and subjects were recruited for the experiment using ORSEE (Greiner, 2015). The points that subjects earned were paid out 1:1 in Euro. In all treatments subjects received

⁶The translated instructions for PD High are provided in Appendix D as an example.

⁷Matching group size can, of course, impact on the ease of coordination and larger matching groups could have made coordination more difficult.

	L	R	Σ
U	0.51	0.22	0.73
D	0.22	0.05	0.27
Σ	0.73	0.27	1

PD High

	L	R	Σ
U	0.61	0.18	0.79
D	0.16	0.05	0.21
Σ	0.77	0.23	1

PD Low

	L	R	Σ
U	0.92	0.04	0.96
D	0.04	0.01	0.04
Σ	0.96	0.04	1

COORD

Note: The sums on the margins show the frequency of the implemented actions. The inner cells show the frequency of outcomes.

Table 3: Actions and Outcomes in Symmetric Treatments

a participation fee of 5 Euro which was paid out regardless of their choices in the experiment. On top of the participation fee subjects earned 13.5 Euro on average. At the end of the experiment subjects were paid out in cash in private. 87 % of subjects were German nationals, 58 % of subjects were male and average age was 24.3. Subjects had been studying for 5.1 semesters on average and they reported to study a wide range of subjects with 8 % studying economics, with 75 % of subjects having had at least one math module during their studies.

3.3 Results

3.3.1 Aggregate Behavior

Table 3 summarizes the implemented actions and outcomes in the form of contingency tables.⁸ As the margins show, U and L are predominantly implemented in all three treatments. In PD High they are implemented 73 % of the time each. This number increases to 79 % respectively 77 % in PD Low and reaches 96 % each in COORD. In order to test if U and L are the modal implemented actions (hypothesis 1), we conservatively consider each matching group an independent observation. If in a matching group U or L is the modal implemented action, we count it as a success. In every treatment U and L are the modal implemented actions in *all* matching groups, yielding 6 successes out of 6 tries for each treatment, generating a p-value of $p = 0.031$ (binomial tests).

Result 1. *In treatments PD High, PD Low and COORD U and L are the modal implemented actions, lending support to hypothesis 1.*

Correspondingly, outcome UL is implemented most of the time in all treatments. The outcome share of UL is 51 % in PD High, rises to 61 % in PD Low and goes as high as 92 % in COORD. In the PD treatments there are still some deviations from UL. Subjects post offers to agents such that the agents deviate from the principals' joint profit maximum UL to UR or DL 22 % of time each. This number decreases to 18 % respectively 16 % in PD Low and finally to 4 % each in COORD. The outcome DR accounts only for 5 % of all outcomes in PD High and PD Low. In COORD outcome DR only accounts for 1 % of all cases. This is important for overall efficiency: Including the welfare of agents, efficiency levels⁹ are 81 % in PD High,

⁸The two trial periods are included in the graphs but omitted from the tables and statistical analyses.

⁹Efficiency = $\frac{\text{achieved payoffs} - \text{minimal payoffs}}{\text{maximal payoffs} - \text{minimal payoffs}}$.


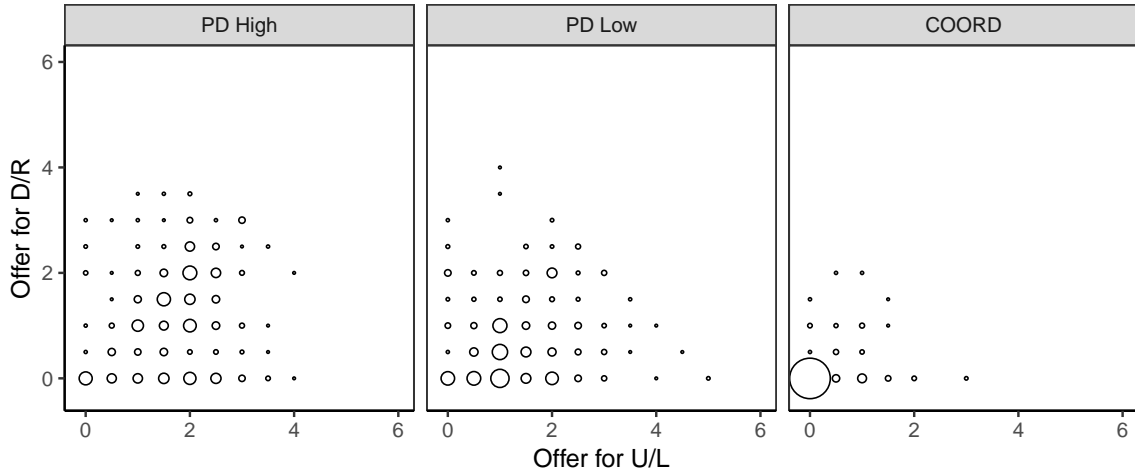


Figure 1: Median Offers to Agents

78 % in PD Low and 92 % in COORD.

Next we turn to the offers that subjects make to agents. First recall the predictions. Hypothesis 2 states that offers decrease from PD High to PD Low to COORD. Indeed, figure 1 shows that median offers for both U/L and D/R decrease from PD High to PD Low to COORD in all incentivized periods. The median offer for $(U/L, D/R)$ decreases from $(1.80, 1.02)$ in PD High to $(1.01, 0.49)$ in PD Low to $(0.00, 0.00)$ in COORD. These offers are significantly different between PD High and PD Low (MWU test, $p = 0.093$ and $p = 0.026$ respectively)¹⁰. Analogously, the offers are significantly different when comparing either PD High or PD Low with COORD ($p = 0.002$ for both auctions). All these results are in line with the theory's comparative statics.

¹⁰Note that in each of the 6 matching groups per treatment we observe 8 subjects making 28 repeated and incentivized decisions. In order to provide conservative estimates we consider the matching group as an independent observation, that is, we average observations on a matching group level before running the tests in this section. All tests in this article are two-sided unless explicitly stated otherwise.



Note: The area of a bubble corresponds to the probability of the corresponding combination of offers.

Figure 2: A Bubble Plot of Transfer Offers

Result 2. *Median offers decrease from PD High to PD Low to COORD, in line with hypothesis 2.*

Condition (AI) predicts symmetric offers with ties that are broken in favor of U and L. Only 2.8 % of offers are tied PD High, 2.7 % are tied in PD Low, whereas in COORD 67.7 % of offers are tied. Out of those 67.7 %, 98.8 % are on equilibrium though. There are very few tied offers off equilibrium. In PD High 2.7 % of offers tie off equilibrium, in PD Low 2.5 % do and in COORD 0.8 % do. Figure 1 already shows that median offers for U/L are remarkably close to the predicted offers of the lowest-transfer equilibria in all three treatments. But note how in all treatments the median offers for D/R are consistently lower or equal to offers for U/L and thus fall short of the symmetry prediction in PD High and PD Low. In all treatments the median offers for U/L are weakly larger than for D/R. In treatment COORD the median offer to either agent is zero, but as the graph shows in the beginning subjects still offer more for U/L than D/R. While figure 1 shows U/L and D/R independently,

we plot the offers in a bubble plot in figure 2.¹¹ Indeed, for PD High and PD Low the offers are predominantly on or below the 45 degree line, i.e. offers are higher for U/L than for D/R. In COORD though, nearly all actions are at or extremely close to zero. In other words, the symmetry prediction seems to work for COORD but not so much for the PD treatments. The difference between offers for D/R and U/L is significant for all treatments (Wilcoxon Signed-Rank, $p = 0.031$ for PD High and PD Low, and, surprisingly, also $p = 0.063$ for COORD). Thus, we have to reject the prediction from hypothesis 3. We will investigate this phenomenon more closely in section 3.4 where we study quantal response equilibria in our games.

Result 3. *Median offers for U/L are larger than for D/R. Offers for U/L follow theoretical predictions but offers for D/R fall short. We reject hypothesis 3.*

Finally, figure 1 indicates that offers decrease weakly over time in all treatments. We test if offers decrease from the first to the second half of the experiment. For PD High the difference is significant for offers for U/L (Wilcoxon Signed-Rank, $p = 0.094$) but not for offers for D/R ($p = 0.156$). For PD Low and COORD the difference is significant for both U/L and D/R ($p = 0.031$ each). While offers decrease a little outcomes are remarkably stable. Table 12 in appendix C shows that the distribution of outcomes does not change much from the first to the second half of the experiment, while offers decrease somewhat.

3.3.2 Individual Learning

Given the complex nature of equilibrium, we were surprised to see aggregate behavior approaching equilibrium play so closely. Of course, the failure of the symmetry

¹¹The distribution does not change much if we look at offers in the second half of the experiment only. Figure 13 in appendix C shows the plot when we only consider offers in the second half of the experiment.

prediction tells us that subjects do not really acquire equilibrium reasoning. Rather, we must look for a combination of initial heuristics and a learning process if we want to understand how behavior develops.

Regarding initial heuristics let us inspect figure 1 again. In all three treatments subjects start out with considerably higher offers for U/L than for D/R. With a little introspection this is perhaps not very surprising. In all three games, implementing U is more important for A principals than implementing D. For example, in PD High, the gain from U is at least 4 while the gain from D is at most 2. So, focusing on the more important auction appears natural. (Indeed, we will revisit this issue—that making a mistake on U is much worse than making a mistake on D—when we estimate quantal response equilibria.) Figure 1 also indicates a potential negative time trend for some offers while in the second half offers seem to stabilize. Indeed, in the first half of the experiment there is a significantly negative time trend for offers for U/L in all three symmetric treatments as well as for offers for D/R in COORD. In the second half, though, there is no significant time trend in any of the treatments. Tables 16 and 17 in appendix C summarize the relevant regression results.

In this section we will focus on the learning process that unfolds in the first half of the experiment without touching on the asymmetry in offers that the principals settle on in the second half of the experiment. We take our cue from Selten’s learning direction theory, that is, the idea that subjects move towards myopic better replies (see for a similar application to learning in auctions Selten et al. (2005)). In the context of our intertwined auctions, the theory simply predicts that a subject will (i) lower her offer if she either won an auction and could have won with a lower offer or if she had offered more than she subsequently gained; and (ii) increase her offer if she lost an auction but could have made a profit by winning it. This leaves two cases. First, a subject may have profitably won an auction by marginally (that is, optimally)

Change	-	NA	+	Σ
	+	81	4	95
0	166	11	58	235
-	198	1	10	209
Σ	445	16	163	624

U/L PD High

Change	-	NA	+	Σ
	+	14	106	43
0	58	170	70	298
-	89	55	19	163
Σ	161	331	132	624

D/R PD High

Change	-	NA	+	Σ
	+	61	3	88
0	169	10	47	226
-	233	2	11	246
Σ	463	15	146	624

U/L PD Low

Change	-	NA	+	Σ
	+	21	67	20
0	63	273	54	390
-	62	63	1	126
Σ	146	403	75	624

D/R PD Low

Change	-	NA	+	Σ
	+	13	7	12
0	116	320	22	458
-	130	4	0	134
Σ	259	331	34	624

U/L COORD

Change	-	NA	+	Σ
	+	6	6	3
0	10	533	25	568
-	26	15	0	41
Σ	42	554	28	624

D/R COORD

Note: On the x-axis LDT shows learning direction theory's predictions. The theory either prescribes a lower offer (-), a higher offer (+) or doesn't yield a prediction (NA). On the y-axis Change shows whether a subject increased (+), decreased (-) or didn't change her offer (0) in the following period. Each cell counts the number of cases. Cells are colored green (dark gray when printed in black and white) if the theory correctly predicted the change, yellow (light gray when printed in black and white) if not, and white if it didn't yield a prediction. Here, we count zeros as successes. Note that we do not use the data from the non-incentivized trial periods.

Table 4: Learning Direction Theory in Periods 1-14

offering more than her opponent. Second, a subject may have lost an auction and could not have profitably deviated because her opponent posted an offer above her reservation value. Any subsequent offer that is lower or equal (or potentially a bit higher) is then a myopic best reply. In these two cases the theory is silent on the direction of an adjustment.

In order to check for directional learning in this spirit we conduct a simple counting exercise. We count how often subjects changed their offer from one period to the next, and how often their behavior is in line with the predictions of learning direction theory. Table 4 shows panels with contingency tables for the first half for all treatments, because this is when we observe a significant time trend in our data. We consider a weak version of learning direction theory and only consider change *against* the prescribed direction as a violation of learning direction theory, counting zero change as a weak success. Cells count the number of cases and are colored green¹² if learning direction theory successfully predicted the change and yellow¹³ if not. The theory turns out to be a good predictor for subjects' behavior if one compares how often subjects move into the predicted rather than the opposite direction. Subjects often stay, but mostly go in the predicted direction. They very seldom go into the opposite direction of what learning direction theory would predict, and this holds for all treatments. In PD High 86 % of all adjustments follow the predictions of learning direction theory, which increases to 89 % in PD Low and 95 % in COORD. Also, note how learning direction theory typically gives clear predictions in PD High and PD Low for offers on U/L. Only 16 respectively 15 observations out of 624 do not come with a prediction. On the other hand, for offers on D/R this number increases to 331 respectively 403 out of 624 observations. This pattern stems from subjects

¹²Dark gray when printed in black and white.

¹³Light gray when printed in black and white.

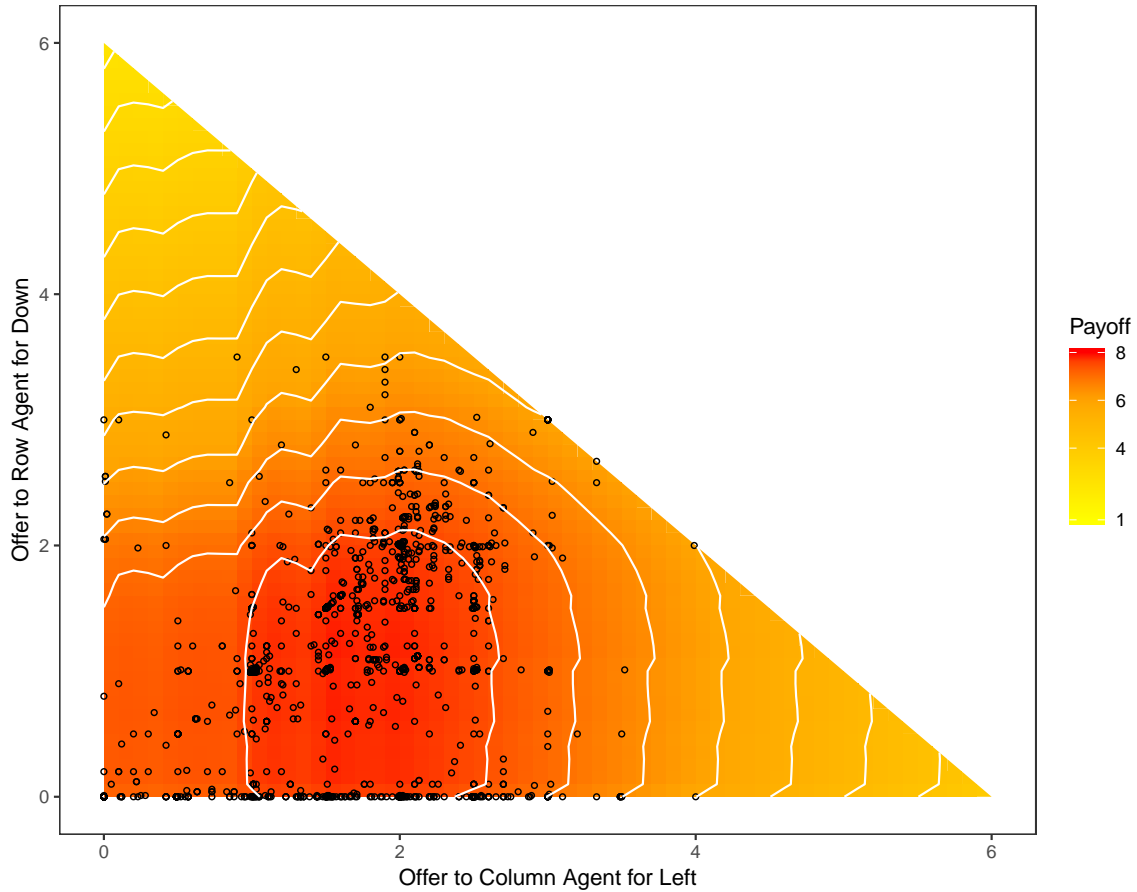
tending to lose the auction for the D/R action so decisively that they could not have gained from either increasing or decreasing their offer. This finding suggests an intuitive explanation for why offers for D/R remain lower than offers for U/L. As subjects could not have gained anyway, there is no force driving them to increase their offers.¹⁴

3.4 Quantal Response Equilibrium

The theoretical predictions for our three symmetric games work well for outcomes but offers on D/R are well below predictions in the two PD treatments. Remember that the lowest-transfer equilibrium offers for $(U/L, D/R)$ in PD High, PD Low and COORD are offers of $(2, 2)$, $(1, 1)$ and $(0, 0)$ respectively. The median offers for U/L are close to the equilibrium predictions but offers for D/R are significantly smaller than the former. Only in COORD the equilibrium predictions for offers fit the data precisely. As discussed above, this appears to stem from some simple heuristic reasoning that encourages principals to offer more in the more important auction—a pattern that is maintained also in the presence of learning (simply because there is not much pressure to increase the offer in an auction one is supposed to lose anyway). In other words, deviations appear to occur when they are not so costly. In order to organize this pattern—that subjects make more errors when they are less costly—estimating a quantal response equilibrium model appears the natural way forward.

Consider a principal A in PD High who expects principal B to choose the symmetric equilibrium offer of $(2, 2)$. If A also offers 2 to both agents, this will implement outcome UL. Principal A ends up only paying the column agent, losing the auction for the row agent, so any potential downward deviations are costless in equilibrium. On the other hand, deviations from the offer of 2 to the column agent are strictly

¹⁴In the second half of the experiment learning direction theory loses some of its bite as subjects tend to stay more than before. Results are provided in table 15 in appendix C.



Note: Each circle represents one observed offer. Given these offers we calculate the expected payoff of every feasible combination of offers and draw the expected payoffs heatmap. The white lines show contour lines between 2.5 and 7.5 in increments of 0.5.

Figure 3: Expected Payoffs for Principal A in PD High

costly: offering too little means losing the auction for the column agent and offering too much means paying too much. As pointed out before, these considerations are suggestive and models that capture the costs of errors should, hence, organize the data and explain the gap in offers for U/L and D/R. The same reasoning holds for PD Low. Moreover, it also provides an explanation for why offers are so close to equilibrium predictions in COORD as in this treatment *all* deviations from the equilibrium offer of 0 are costly. Still, we fit the QRE to all three treatments.


In figure 3 we draw the expected payoff heatmap of a principal A playing against

the empirical distribution of offers in PD High. It shows an area of approximately maximal payoffs, the darkest part of the figure, around the perpendicular on an offer for (L, D) of $(2, 0)$, bounded above by the equilibrium offer of $(2, 2)$. This corridor happens to contain most actual offers indicating that subjects play noisy best replies against the empirical distribution.

We formalize this intuition in a QRE framework following McKelvey and Palfrey (1995). We do this in order to show that the above intuition can be captured in a consistent equilibrium framework. As mentioned before there is evidence that offers change from the first to the second half of the experiment. Figure 1 and correspondingly table 12 in Appendix C) show a time trend in the first half of the experiment which is explained by directional learning. Regressions show a significant time trend in some of the offers in the first half of the experiment which dies down in the second half (see tables 16 and 17 in appendix C). Consequently, in what follows we consider only periods 15 to 28.¹⁵

Finally, we have to tackle the issue of multiplicity of equilibria. There could be multiple QREs in the games that we are looking at. While in simple games it is feasible to calculate QREs explicitly it is much more difficult in complex games such as GPTAs. Instead, we check the robustness of our results by feeding a large number of different distributions of play into our algorithm to test whether the algorithm converges to different equilibria. We use the empirical distribution of play, a uniform distribution, as well as a large number of distributions covering the simplex in a grid search as starting points and find that the algorithm converges to the same QRE in all treatments. While this does not prove uniqueness there is certainly a very large basin of attraction for the QRE that we compute. More details on the setup of the QRE as well as the robustness checks are documented in Appendix B).

¹⁵See Goeree et al. (2002) for a similar procedure.



Note: The area of a bubble corresponds to the probability of the corresponding combination of offers. Very small bubbles may not be displayed.

Figure 4: Empirical and Best-Fit QRE Distribution of Offers

The best fits are obtained for a precision parameter $\lambda = 1.91$ for PD High, $\lambda = 2.27$ for PD Low and $\lambda = 5.23$ for COORD. The models have a negative log-likelihood of 2466.614 for PD High, 2200.297 for PD Low and 177.4563 for COORD. The QRE models fit the data well. Figure 4 juxtaposes the empirical distribution of offers with the predicted probabilities in the best-fitting QRE. The area of a bubble corresponds to the probability of the corresponding combination of offers. Note how closely the distributions resemble each other. The average absolute difference between the empirical distribution of play and the QRE is just 0.007 in PD High, 0.013 in PD Low and 0.005 in COORD. Empirically, there are a lot of offers different from $(2, 2)$ in PD High with more distribution mass on low offers for D/R than on low offers for U/L. Similarly, in the QRE with the best fit there is a lot of probability mass on low offers for D/R but less so for offers for U/L. These results are in line with our earlier intuition. Errors for offering too little for D/R are not very costly given the other subject's offer but conversely for U/L they are. A similar pattern can be observed in PD Low but now probability mass is shifted to lower offers in both dimensions, which is in line with theoretical predictions. Finally, for COORD nearly all observations are at or extremely close to $(0, 0)$ which is also reflected in the corresponding QRE.

Result 4. *Subjects' overall choice patterns are consistent with subjects making errors in a quantal response equilibrium framework.*

In this section we have shown how quantal response equilibrium can make sense of the behavioral patterns in the data. But, obviously, we can explain a lot of behavior when fitting a model to data ex post. Consequently, we designed an asymmetric follow-up treatment to test some predictions of QRE out of sample.

4 An Asymmetric Game

4.1 Treatment and Hypotheses

All games so far were symmetric. In this section we study an asymmetric game to test predictions of quantal response equilibrium out of sample.¹⁶ Testing the QRE model is a challenge. To formulate a hypothesis we have to predict behavior, so one might be inclined to estimate a joint precision parameter across the symmetric treatments and use it to predict a distribution of play in the new treatment. But comparing precision parameters across games is notoriously difficult. First, games may differ in complexity and as such subjects may make errors at a different rate – evidence for which can be gleaned from the different precision parameters that we estimate in the symmetric treatments. The COORD game is arguably a simpler game than the PD games and as such yields the highest precision parameter in the QRE estimation. Second, the precision parameter is sensitive to payoffs. A simple mathematical operation such as scaling up payoffs yields different parameter estimates. Hence, fitting QREs necessarily comes with a substantial degree of freedom. Instead of predicting and testing against specific distributions of play, we use QRE to predict offer patterns that hold true for *any* precision parameter but are not predicted by subgame perfect Nash equilibrium.

The asymmetric treatment features a gross payoff matrix corresponding to a simple dominance solvable game and we refer to this treatment as DOM. In the symmetric PD treatments we observe an asymmetry within a principal's offers as principals typically offer more for L than for D. In DOM we predict instead an asymmetry of offers *between* the two types of principals. The payoff matrix is shown in table 5.

¹⁶We also implemented a second asymmetric treatment with a gross payoff matrix corresponding to a battle of the sexes game to test if the theory has bite in a more intricate setting. Results are reported in AppendixA.

	L	R
U	5, 0	0, 1
D	0, 1	0, 2

Table 5: DOM Gross Principals' Payoff Matrices Depending on Agents' Actions




Figure 5: Predicted Average QRE Offers in DOM

Principal A offers transfers for actions U and L while principal B offers transfers for D and R. Outcome UL is the most efficient so theory predicts it to arise in equilibrium. The corresponding offers that support this outcome are all equal to 1. As before, subjects are randomly assigned to matching groups of eight (with four principals A and B each) and are rematched every round, playing again for 2 trial and 28 incentivized rounds. In DOM we ran two sessions of 24 subjects each.¹⁷

Let us consider what QRE predicts in this treatment. Principal A stands to lose a lot from losing either U or L so she should make few errors. Principal B on the other hand can expect to lose both actions D and R so she is expected to randomize

¹⁷Calculations are documented in appendix C and more details to the experimental implementation of the asymmetric treatments are documented in appendix A.

between losing bids. Figure 5 shows how QRE’s predicted offers for both principals change as the precision parameter λ increases.¹⁸ At $\lambda = 0$ both principals perfectly randomize between all available offers. As λ increases average offers decrease. As $\lambda \rightarrow \infty$ principal A’s offers converge to 1 (the subgame perfect Nash equilibrium) for both U and L. Principal B on the other hand expects to lose against principal A’s offers and is indifferent between all offers below 1. Hence, as $\lambda \rightarrow \infty$ principal B’s average offers converge to 0.5 for both D and R. From this graph it is easy to see that for any $\lambda > 0$ QRE predicts principal A’s offers to be larger than principal B’s offers while standard theory predicts all offers to be the same. Furthermore, note that since the offers for U and L are bigger than the offers for D and R, both QRE and standard theory predict that U and L should typically be implemented.

Hypothesis 4. *U and L are the modal implemented actions in DOM.*

Hypothesis 5. *All offers tie in DOM. (Standard theory’s null hypothesis)*

Hypothesis 6. *Offers for D respectively R are smaller than offers for U respectively L in DOM. (QRE’s alternative hypothesis)*

4.2 Results

As in the symmetric treatments, Prat and Rustichini’s (2003) theory does a remarkable job predicting the implemented actions and resulting outcomes as table 6 shows. In treatment DOM, U and L are the implemented actions 83 % respectively 80 % of the time. In all 6 matching groups U and L are the modal implemented actions, providing evidence for hypothesis 4 (two-sided binomial tests, $p = 0.016$). Corresponding to the implemented actions, in DOM 72 % of outcomes are UL. Only 19 %

¹⁸Note that, on average, principal A’s offers for U and L should be the same as well as principal B’s offers for D and R.

	L	R	Σ
U	0.72	0.11	0.83
D	0.08	0.09	0.17
Σ	0.80	0.20	1

Note: The sums on the margins show the frequency of the implemented actions. The inner cells show the frequency of outcomes.

Table 6: Actions and Outcomes in DOM




Figure 6: Median Offers in DOM

of outcomes are either DL or UR and only 9 % of the time principal B can implement her most preferred outcome DR. As even minor deviations from UL are very costly from an efficiency standpoint subjects achieve an efficiency level of only 74 %.

Result 5. *In DOM U and L are the modal implemented actions, supporting hypothesis 4.*

Let us now turn to the offers. Nash equilibrium would predict all offers to tie. Figure 6 shows that this is not the case. Principal A's median offers for (U, L) are $(1.05, 1.10)$ while principal B's median offers for (D, R) are $(0.70, 0.75)$. So, on

the one hand, principal A's offers for U and L are not significantly different from each other as are principal B's offers for D and R (Wilcoxon Rank-sum tests, $p = 1.000$). But on the other hand, principal A's offers for U and L are both significantly larger than principal B's offers for D and R (Wilcoxon Rank-sum tests, $p = 0.031$). Hence, we can reject standard theory's null hypothesis 5 but we cannot reject QRE's alternative hypothesis 6. Examining the offers more closely, principals show a taste for symmetry. Principal A offers symmetric transfers to both agents 73 % of the time while principal B offers a symmetric transfer 51 % of the time. Indeed, in a QRE framework it makes sense for principal A to be more interested in offering a symmetric transfer than principal B. Principal B typically loses against principal A, so, in terms of payoffs, it makes little difference if she is to lose with an asymmetric or a symmetric bid. For principal A, on the other hand, winning both auctions is equally important, hence we should intuitively expect more symmetric offers by principal A than by principal B.

Result 6. *In DOM offers for U and L are larger than offers for D and R. Hence, we have to reject standard theory's null hypothesis 5 in favor of QRE's alternative hypothesis 6.*

In DOM the best-fit QRE has a precision parameter of $\lambda = 3.07$ and the model has a negative log-likelihood of 3668.179.¹⁹ Figure 7 shows the offers that we observe in the experiment and the offers of the best-fit QRE. Note how closely the distributions resemble each other. QRE predicts principal B to randomize between losing offers,

¹⁹Subjects do not seem to engage much in learning as offers do not significantly increase from the first half to the second half of the experiment (Wilcoxon Rank-sum tests, $p = 0.313$ and $p = 1.000$ for principal A's and $p = 0.8438$ for both of principal B's offers). Medians are also stable over time as table 13 in Appendix C shows. Similarly, a linear regression of offers on periods shows no significant time trend. Regression results are reported in Table 19 in Appendix C. Hence, we fit QRE on the data of all 28 periods. In the asymmetric game there are two types of principals instead of just one so we adapt our estimation strategy accordingly. Details and robustness checks are documented in Appendix C.

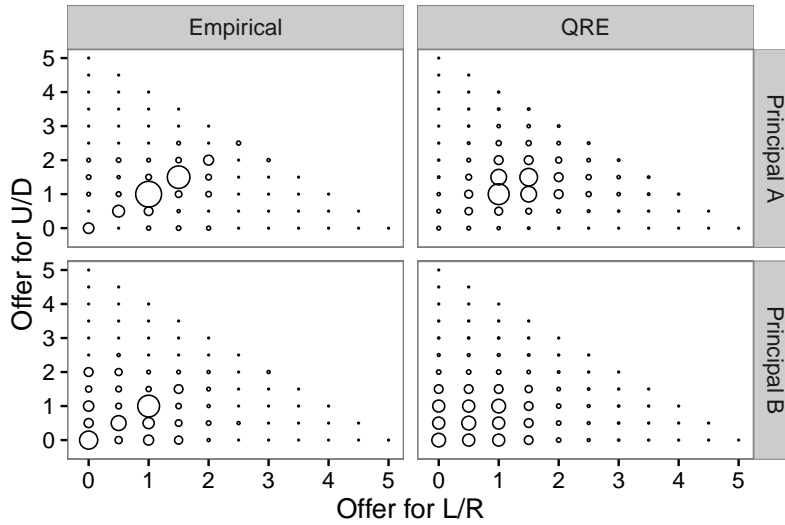


Figure 7: Empirical and Best-Fit QRE Distributions for DOM

hence principal B offers should be concentrated in the lower left corner of the figure. Principal A, on the other hand, is expected to mostly offer symmetric offers of (1,1) and slightly larger.²⁰ While we observe slightly more symmetric offers in the data than we would expect from QRE, the general pattern is borne out. A substantial share of principal B's offers are asymmetric and principal A's offers are mostly the symmetric offers of (1,1) and (1.5, 1.5). The mean absolute difference between the empirical distribution of play and the QRE is 0.01.

Result 7. *In DOM quantal response equilibrium organizes subjects' offers very well.*

5 Discussion

This paper presents an experimental test of Prat and Rustichini (2003). Examining situations where multiple principals can influence multiple agents by promising monetary transfers, their model captures an important class of real-world strategic

²⁰This pattern can also be seen in Figure 15 in Appendix C, which shows a number of QREs for increasing λ s.

interactions. Yet, so far the model has remained empirically untested. Despite its non-trivial equilibrium logic, the model predicts behavior well. In our symmetric treatments as well as in the asymmetric DOM treatment the theory successfully predicts which actions are implemented about 80 % of the time.

While our results for implemented actions provide strong support for the theory, subjects' offers require closer scrutiny. Comparative static predictions for the symmetric games that we implemented hold. Also, behavior quickly settles down as subjects learn to play the game. The observed learning pattern is well organized by Selten's learning direction theory. While the naive prediction of all probability mass being concentrated on one point fails, deviations arise in a clear pattern. On the one hand, offers for actions that are implemented (such that one has to pay for them) are very close to the theory's predictions. On the other hand, offers for actions that are typically not implemented (and thus are without immediate payoff consequences) are lower than predicted. Instead of making offers for the latter such that their opponent becomes indifferent subjects bid too little violating the agent indifference condition of Prat and Rustichini's (2003) theory. To explain the deviations we apply a quantal response equilibrium framework to our data. QRE can make sense of the asymmetry in offers that we observe and allows us to successfully predict behavior out of sample in a follow-up treatment. In treatment DOM we test for an asymmetry between two types of principals' offers as a central *ex ante* prediction of QRE. We can reject symmetry of offers in favor of QRE's predictions. Integrating QRE into GPTAs is natural due to the inherent asymmetry in local incentives at the equilibrium. Finally, it should be noted that, on the level of implemented outcomes and actions, QREs are, for the usual range of the precision parameter, remarkably close to Nash.

One might be tempted to compare our data to behavior in simple 2×2 games played without intermediaries. We do not offer such a comparison as it is misleading.

The whole idea of GPTAs is to look at situations where principals *cannot* play the base game (just as lobbyists cannot vote in parliament). Also, principals do not simply hire an agent to play a game for them. Rather, they influence other people who play a game whose outcome has externalities—on the principals. Moreover, the game that these others play, the agents in the GPTA terminology, is in our case a 2×2 game where all players receive zero payoffs in all cells, that is, the base game is neither a prisoners’ dilemma, nor a coordination game. Instead, consider that from the principals’ perspective GPTAs have a similar structure as two simultaneous first-price sealed-bid auctions with synergies, which have been studied extensively. Synergies between auctions give rise to the threat of an inefficient allocation of items. Rassenti et al. (1982) study combinatorial auctions as a solution to this issue and apply it to the allocation of airport slots. Another application is spectrum auctions, in which telecommunication firms bid for broadcasting rights (see, for example, McAfee and McMillan (1996) for an analysis of the FCC spectrum auctions in the 1990s). An experiment that is related to our setup is Chernomaz and Levin (2012), who compare bidding in simultaneous first-price sealed-bid auctions where bidding on packages of items is either allowed or not, arguing that package bidding raises efficiency. For further reading, see reviews on multiunit auctions in Kwasnica and Sherstyuk (2013) and Kagel and Levin (2015). However, despite their similarities with auctions GPTAs have some important differences that make it difficult to compare our results directly to the results in the literature on multiunit auctions. Like bidders in an auction, principals in a GPTA have (heterogeneous) valuations of agents’ actions and post sealed offers to agents. But, typically, bidders’ valuations in an auction are private information while in a GPTA the principals’ valuations are public information. Hence, offering transfers in GPTAs is perhaps more comparable to posting prices in a Bertrand game.

There are some experimental auction papers that investigate settings in which valuations are common knowledge. However, these papers only look at single auctions. Andreoni et al. (2007) run a laboratory study on the effect of information about other bidders' valuations on bidding behavior in first- and second-price sealed-bid auctions. In one of their treatments (*Auction 3*) bidders have full information about the other bidders' valuations. This transforms the auction into a Bertrand game with asymmetric production costs. They find that, as theory predicts, the subject with the highest valuation typically submits a bid close to the second-highest valuation, while the subject with the second-highest valuation bids either her valuation or below. This behavioral pattern is very similar to our findings. Boone et al. (2012) find similar results studying straightforward Bertrand games with asymmetric production costs in the laboratory. Dugar and Mitra (2016) investigate a similar setting but vary the cost differences. They find that prices increase when the absolute difference between the two marginal costs decreases.

Our study should be seen as a first step towards more research into GPTAs. We only test relatively simple interactions (2 principals and 2 (computerized) agents with 2 strategies each). Including human agents appears to be one of the more desirable extensions for future work. Of course, this route is likely to require the consideration of social preferences which are absent from the base model. (On the other hand, the uncertainty that the human factor would introduce would render these GPTAs closer to auctions with incomplete information.) Other obvious extensions would examine more complicated types of interaction. How far can the theory be pushed in explaining who gets their way in such games of economic influence? Voting games appear to be a particularly attractive avenue given the model's natural applications to lobbying. Our results suggest that, just as theory predicts, principals tend to incentivize agents such that the socially efficient outcome results most of the time. It

would be intriguing to see similar results in voting games.

References

Aidt, Toke S. and Uk Hwang (2008) “On the Internalization of Cross-National Externalities through Political Markets: The Case of Labour Standards,” *Journal of Institutional and Theoretical Economics*, Vol. 164, pp. 509–533, URL: <http://openurl.ingenta.com/content/xref?genre=article&issn=0932-4569&volume=164&issue=3&spage=509>, DOI: <http://dx.doi.org/10.1628/093245608785363380>.

Andreoni, James, Yeon-Koo Che, and Jinwoo Kim (2007) “Asymmetric information about rivals’ types in standard auctions: An experiment,” *Games and Economic Behavior*, Vol. 59, pp. 240–259, DOI: <http://dx.doi.org/10.1016/j.geb.2006.09.003>.

Bikhchandani, Sushil (1999) “Auctions of Heterogeneous Objects,” *Games and Economic Behavior*, Vol. 26, pp. 193–220, DOI: <http://dx.doi.org/10.1006/game.1998.0659>.

Boone, Jan, María Jose Larraín Aylwin, Wieland Müller, and Amrita Ray Chaudhuri (2012) “Bertrand competition with asymmetric costs: Experimental evidence,” *Economics Letters*, Vol. 117, pp. 134–137, URL: <https://linkinghub.elsevier.com/retrieve/pii/S016517651200239X>, DOI: <http://dx.doi.org/10.1016/j.econlet.2012.04.098>.

Bordignon, Massimo, Luca Colombo, and Umberto Galmarini (2008) “Fiscal federalism and lobbying,” *Journal of Public Economics*, Vol. 92, pp. 2288–2301, URL:

<http://linkinghub.elsevier.com/retrieve/pii/S0047272708000911>, DOI:
<http://dx.doi.org/10.1016/j.jpubeco.2008.05.003>.

Chernomaz, Kirill and Dan Levin (2012) “Efficiency and synergy in a multi-unit auction with and without package bidding: An experimental study,” *Games and Economic Behavior*, Vol. 76, pp. 611–635, URL: <https://linkinghub.elsevier.com/retrieve/pii/S0899825612001145>, DOI: <http://dx.doi.org/10.1016/j.geb.2012.07.008>.

Dugar, Subhasish and Arnab Mitra (2016) “Bertrand Competition with Asymmetric Marginal Costs,” *Economic Inquiry*, Vol. 54, pp. 1631–1647, URL: <http://doi.wiley.com/10.1111/ecin.12328>, DOI: <http://dx.doi.org/10.1111/ecin.12328>.

Fischbacher, Urs (2007) “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental Economics*, Vol. 10, pp. 171–178, URL: <http://link.springer.com/10.1007/s10683-006-9159-4>, DOI: <http://dx.doi.org/10.1007/s10683-006-9159-4>.

Fredriksson, Per G. and Daniel L. Millimet (2007) “Legislative Organization and Pollution Taxation,” *Public Choice*, Vol. 131, pp. 217–242, URL: <http://link.springer.com/10.1007/s11127-006-9114-0>, DOI: <http://dx.doi.org/10.1007/s11127-006-9114-0>.

Goeree, Jacob K., Charles A. Holt, and Thomas R. Palfrey (2002) “Quantal Response Equilibrium and Overbidding in Private-Value Auctions,” *Journal of Economic Theory*, Vol. 104, pp. 247–272, URL: <http://linkinghub.elsevier.com/retrieve/pii/S002205310192914X>, DOI: <http://dx.doi.org/10.1006/jeth.2001.2914>.

- Greiner, Ben (2015) "Subject pool recruitment procedures: organizing experiments with ORSEE," *Journal of the Economic Science Association*, Vol. 1, pp. 114–125, URL: <http://link.springer.com/10.1007/s40881-015-0004-4>, DOI: <http://dx.doi.org/10.1007/s40881-015-0004-4>.
- Judd, Kenneth L. (1998) *Numerical Methods in Economics*, Cambridge, MA: MIT Press.
- Kagel, John H. and Dan Levin (2015) "Auctions: A Survey of Experimental Research," in John H. Kagel and Alvin E. Roth eds. *The Handbook of Experimental Economics*, Vol. 2, Princeton, New Jersey: Princeton University Press, pp. 290–346.
- Kwasnica, Anthony M. and Katerina Sherstyuk (2013) "Multiunit Auctions," *Journal of Economic Surveys*, Vol. 27, pp. 461–490, URL: <http://doi.wiley.com/10.1111/joes.12017>, DOI: <http://dx.doi.org/10.1111/joes.12017>.
- McAfee, R. Preston and John McMillan (1996) "Analyzing the Airwaves Auction," *Journal of Economic Perspectives*, Vol. 10, pp. 159–175, URL: <http://pubs.aeaweb.org/doi/10.1257/jep.10.1.159>, DOI: <http://dx.doi.org/10.1257/jep.10.1.159>.
- McKelvey, Richard D. and Thomas R. Palfrey (1995) "Quantal Response Equilibria for Normal Form Games," *Games and Economic Behavior*, Vol. 10, pp. 6–38, DOI: <http://dx.doi.org/10.1006/game.1995.1023>.
- Prat, Andrea and Aldo Rustichini (2003) "Games Played Through Agents," *Econometrica*, Vol. 71, pp. 989–1026, URL: <http://doi.wiley.com/10.1111/1468-0262.00437>, DOI: <http://dx.doi.org/10.1111/1468-0262.00437>.

Rassenti, Stephen J., Vernon L. Smith, and Robert L. Bulfin (1982) “A Combinatorial Auction Mechanism for Airport Time Slot Allocation,” *The Bell Journal of Economics*, Vol. 13, pp. 402–417, DOI: <http://dx.doi.org/10.2307/3003463>.

Selten, Reinhard, Klaus Abbink, and Ricarda Cox (2005) “Learning Direction Theory and the Winner’s Curse,” *Experimental Economics*, Vol. 8, pp. 5–20, URL: <http://link.springer.com/10.1007/s10683-005-1407-5>, DOI: <http://dx.doi.org/10.1007/s10683-005-1407-5>.

	L	R
U	5, 2	0, 0
D	0, 0	2, 4

Table 7: BoS Gross Principals' Payoff Matrices Depending on Agents' Actions

Appendix A Asymmetric Treatments

The experimental procedures were nearly the same as in the symmetric treatments. The main difference is that we now also employed different instructions such that the principals would see the game from their own perspective. Subjects are assigned the role of principal A or principal B at the beginning of the game and they keep their role throughout the experiment. In all asymmetric treatments, principal A can incentivize the column agent to choose L and the row agent to choose U. Conversely, principal B can incentivize the column agent to choose R and the row agent to choose D. Subjects in both treatments were endowed with a budget of 5.

The second, more intricate, asymmetric treatment has a gross payoff matrix corresponding to a battle of the sexes. We refer to this treatment as BoS. The payoff matrix is shown in table 7. In total we ran four sessions.²¹

The unique efficient outcome in this game is UL, hence it must result in any pure-strategy Nash equilibrium. Implementing outcome DR instead would earn principal B an increase in gross payoffs of 2. Hence, in the lowest-transfer equilibrium principal A must offer a total transfer of 2 to both agents and principal B must match these offers. For the BoS game the predictions are not as crisp as for the other treatments. Before, we had a unique lowest-transfer equilibrium and a point prediction for the individual offers. Now we have a point prediction for the total offer but not for the individual

²¹Three sessions were run with 24 subjects, but in a fourth session too few subjects showed up and the session was run with 20 subjects, with one matching group of eight and two matching groups of six. In our analysis we drop the entire session, but all results are robust to including these observations.

	L	R	Σ
U	0.59	0.02	0.61
D	0.02	0.37	0.39
Σ	0.61	0.39	1

Note: The sums on the margins show the frequency of the implemented actions. The inner cells show the frequency of outcomes.

Table 8: Actions and Outcomes in BoS

offers. Any combination of offers that sum up to 2 can support a lowest-transfer equilibrium. Consequently, the BoS game exhibits two added layers of complexity compared to the symmetric games. First, there is now more conflict from the outset. Second, subjects also have to coordinate on how to allocate the sum of offers. After the strong performance of Prat and Rustichini (2003)'s theory for symmetric games, the BoS game serves very much a stress test. The predictions are

Hypothesis 7. *U and L are the modal implemented actions in BoS.*

and

Hypothesis 8. *Offers to both agents sum to 2 for both principals in BoS.*

U and L are the implemented actions in 61 % of all cases. In 7 out of 9 matching groups U and L are the modal implemented actions, providing evidence for hypothesis 7 (one-sided binomial tests, $p = 0.090$)²². 59 % of outcomes are UL with DR coming up 37 % of the time and remarkably little miscoordination on either UR or DL. These results are interesting from an efficiency point of view as subjects successfully manage to avoid miscoordination. Subjects achieve a remarkable efficiency level of 90 %. The distribution of outcomes is also quite stable in either treatment and does not change much over time as tables 13 and 14 in appendix C show.

²²Note that we construct our tests very conservatively by dropping one session with too few subjects, aggregating on the matching group as the independent observation and making no distributional assumptions. If one includes the omitted session the success ratio for hypothesis 7 increases to 10 out of 12 (two-sided binomial tests, $p = 0.019$).

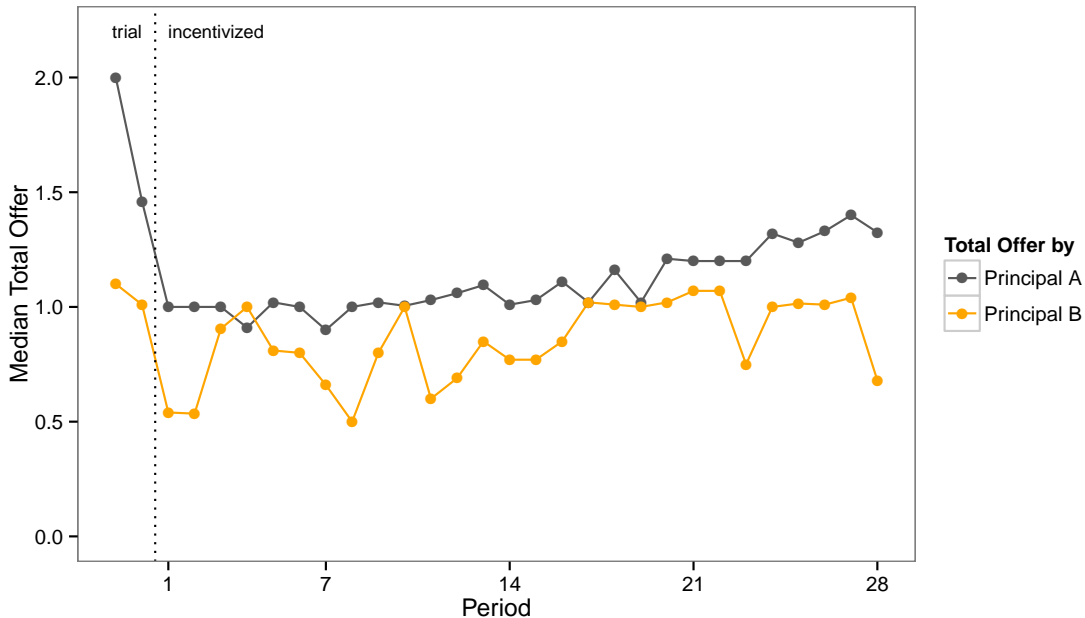


Figure 8: Median Total Offers in BoS

When looking at the offers in BoS let us first check how subjects deal with the issue of coordination. Even more often than in treatment DOM, subjects typically choose the same offers for the two actions. 81 % of the time principals choose to offer the same amount to both agents. Choosing symmetric offers appears to be a focal point on which subjects coordinate. Intuitively, one can see the appeal of such offers. Given that other subjects make symmetric offers it is best to place a symmetric offer as well. Either one wins with both offers and ends up with one's preferred outcome, or one loses and ends up with the opponent's preferred outcome. Either outcome is preferable to miscoordination on UR or DL though, which would be the potential result of posting a non-symmetric offer. As our predictions concern total offers and most offers are symmetric anyway we can focus on total offers. Figure 8 shows the evolution of total offers over the course of the experiment. Total offers to agents seem to increase somewhat but remain below the equilibrium prediction of 2. Furthermore,

principal A's median total offer is 1.08 which is significantly higher than principal B's median total offer of 0.82 (Wilcoxon Rank-sum, $p = 0.098$).²³

Result 8. *In BoS median total offers are below equilibrium predictions and principal A offers more than principal B. Therefore we have to reject hypothesis 8.*

The best fit in the BoS treatment is for $\lambda = 1.62$ and the model has a negative log-likelihood of 4912.353.²⁴ Figure 9 summarizes the results. The QRE does not fit as well as in the case of the symmetric treatments but still captures the essential features of the data. The mean absolute difference between the empirical distribution of play and the QRE is 0.011. As mentioned earlier, subjects predominantly use symmetric offers, but there is no force in the QRE that would ensure such symmetry. Consequently, in the corresponding QRE there is too much probability mass on asymmetric offers. As principal A tends to offer more than principal B, principal B randomizes on low symmetric as well as asymmetric offers. Still, principal A is predicted to offer more than principal B, which is reflected in the data.

Appendix B QRE Procedures

B.1 Symmetric QRE

A principal A chooses an action out of her strategy set $X = \{x_1, \dots, x_J\}$ which consists of J pure strategies. In the symmetric treatments there is no meaningful distinction between principals A and B. For example principal B's offer of $(t_R^B(U), t_C^B(R)) = (1, 0)$ for row and column agent respectively is equivalent to principal A's offer of

²³As before we average observations on a matching group level before running the tests in this section.

²⁴As above, we also calculate a number of QREs for an increasing λ in order to understand how the process converges. Results are provided in figure 16 in appendix C.

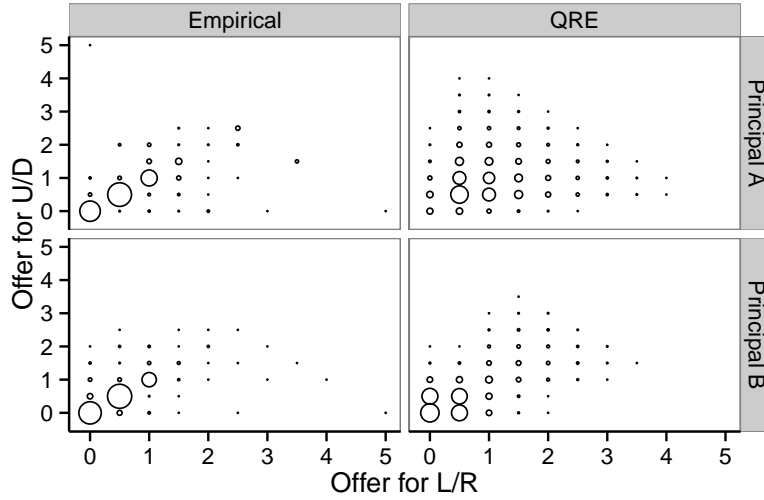


Figure 9: Empirical and Best-Fit QRE Distribution of Offers for BoS

$(t_R^A(D), t_C^B(L)) = (0, 1)$. Furthermore, all subjects perceived the game as playing as principal A. Therefore, we only need to model principal A. Her opponent principal B's strategy is constructed by rearranging the principal A's own strategy. Principal A's strategy is a vector σ_i of length J that assigns each pure strategy a probability and adds up to 1. Given σ_i we can rearrange the pure strategies and their probabilities to construct an equivalent principal B. Given this principal B's strategy and probability vectors we can calculate the expected payoff $E\pi(x)$ for each of principal A's actions. Principals choose actions with a logit choice function and a precision parameter λ . For $\lambda = 0$ a principal randomizes uniformly between all actions but as λ grows errors become smaller and smaller until behavior approaches for $\lambda \rightarrow \infty$ a Nash equilibrium. A quantal response equilibrium for a given λ is characterized as a fixed point of such a logit response function. The probability of choosing an action $x \in X$ is as follows.

$$Prob(x) = \frac{\exp(\lambda E\pi(x))}{\sum_{t \in X} \exp(\lambda E\pi(t))}$$

In the theoretical model the action space for offers is continuous. A logit choice

models needs discrete actions though. Therefore, we approximate the model by discretizing the action space (resembling, of course, the experimental environment). Choosing the appropriate discrete step size is not trivial. The smaller the step size the bigger the computational demand. This issue is exacerbated in our case as we have a two-dimensional action space. The number of distinct actions increases nearly quadratically in the step size.²⁵ Furthermore in order to calculate the expected payoffs each action has to be evaluated against every possible other action. We choose a step size of 0.5 as it computes the equilibria reasonably quickly and as it is sufficient to illustrate our point.²⁶ We define the set of actions as the set of actions that fulfill the budget constraint G . In PD High the budget is 6, in PD Low 5, and in COORD 3.

$$X = \{(0, 0), (0, 0.5), \dots, (0, G), (0.5, 0), (0.5, 0.5), \dots, (G, 0)\}$$

For our analysis we need two algorithms. The first is the equilibrium algorithm which computes a logit equilibrium for a given λ . The second is the optimization algorithm which determines the λ that fits the data best. For a given λ we use a fixed point iteration approach (see, for example, Judd (1998)). Using some probability distribution as a starting point the equilibrium algorithm computes a logit response which is used as an input for another logit response and so on. This algorithm then continues until it converges.²⁷ To check robustness we use a large number of different starting distributions. We use a uniform distribution where a subject randomizes uniformly over all actions, we use the empirical distribution of play, and we use a set

²⁵For a budget G and a step size α subjects have $\frac{1}{2}(\frac{G^2}{\alpha^2} + \frac{G}{\alpha})$ feasible strategies. This leads to $\frac{1}{4}(\frac{G^2}{\alpha^2} + \frac{G}{\alpha})^2$ combinations of strategies one has to consider in order to calculate a logit response. In treatment PD High the budget is $G = 6$. For a step size of $\alpha = 0.5$ we have 6084 comparisons. For a step size of $\alpha = 0.1$ this increases to about 3.5 million comparisons. For a step size of $\alpha = 0.01$ (the actual step size in the experiment) this increases to about 32 billion comparisons.

²⁶The results do not change qualitatively for a smaller step size.

²⁷A sequence never perfectly converges, so we consider a sequence as converged when the sum of the absolute probability differences between the last two iterations is smaller than 0.001.

of 1792 randomly generated distributions. For these randomly generated distributions we generate offers from a two-dimensional normal distribution. Given a budget G each dimension has its own mean $\in (G/16, G/8, G/4, G/2)$ and variance $\in (0.1, 0.4, 0.7, 1)$, and the offers are correlated by $\rho \in (-0.9, -0.6, -0.3, 0, 0.3, 0.6, 0.9)$.

For each of the randomly generated distributions we draw 1000 offers to the row and column agent from. Any offer that is below 0 and any offer for which the sum of offers is above the budget is thrown away. The remaining offers are translated into a distribution of play. The algorithm converges to the same QRE no matter which distribution we use as a starting value in all treatments.

The second algorithm computes the λ and assorted equilibrium that fit the data best. Given the data the algorithm finds an equilibrium that minimizes the negative log-likelihood using the standard R optim function.²⁸ Our empirical data is close to continuity with subjects being able to specify offers in steps of 0.01. In order to fit the steps in our QRE we round each offer to its next QRE step, that is, every offer weakly smaller than 0.25 is counted as 0, every offer larger than 0.25 and weakly smaller than 0.75 is counted as 0.5 and so on. For our fits to make sense we have to assume the observations to be independent and to come from the same distribution.

B.2 Robustness Checks

The QREs are fitted on individual choice data of tuples of offers for $(U/L, D/R)$ so we check if the equilibria are also consistent with other patterns in the data. Figure 10 shows histograms for empirically observed offers and those predicted by the best-fitting QREs separately. The QRE predictions correspond nicely to the observed offers. Recall that we fit the data on a two-dimensional action set, whereas now we

²⁸We also calculate a number of QREs for an increasing λ in order to understand how the process converges. Results are provided in figure 14.

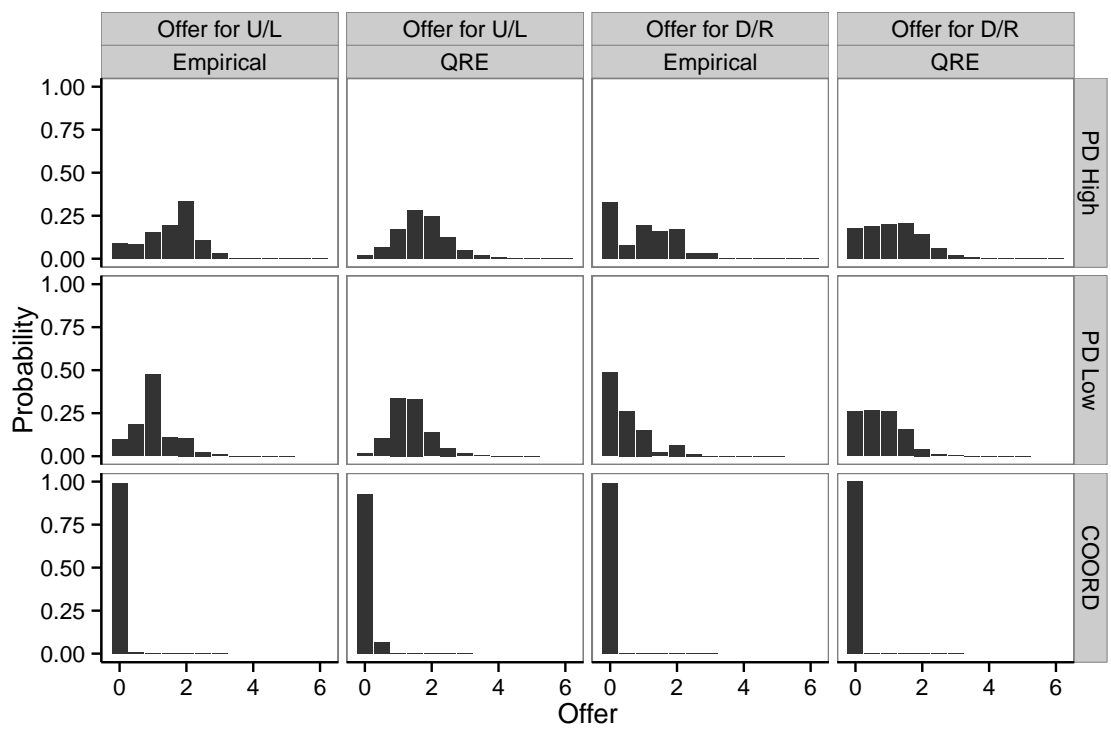


Figure 10: Histograms of Empirical and Best-Fit QRE Offers

		PD High		PD Low		COORD	
		Emp.	QRE	Emp.	QRE	Emp.	QRE
Offers	U/L	1.56	1.71	1.04	1.42	0.01	0.04
	D/R	1.03	1.13	0.49	0.81	0.01	0.00
Outcomes	UL	0.50	0.59	0.63	0.71	0.96	1.00
	UR	0.21	0.18	0.17	0.13	0.01	0.00
	DL	0.25	0.18	0.14	0.13	0.03	0.00
	DR	0.05	0.04	0.05	0.02	0.00	0.00

Table 9: Empirical and Best-Fit QRE Mean Offers and Outcomes

look at offers in one dimension at a time, so such a similarity does not follow by construction. The empirical mean offers are also very close to their corresponding predicted mean offers. Table 9 shows the observed and predicted mean offers. In all treatments predicted offers are reasonably close to the empirical means. The QREs also manage to capture that offers for U/L are higher than for D/R. Below the mean offers the same table also shows the observed and predicted distribution of outcomes. In all QREs outcome UL is predicted to be modal, followed by UR and DL and finally DR, just like we observe empirically. All in all, the QREs consistently fit the empirical data in multiple dimensions.

B.3 Asymmetric QRE

For the asymmetric games we calculate quantal response equilibria in a way similar to the symmetric treatments. The main difference is that due to the asymmetry in payoffs we now have to model two principals with choice probability vectors σ_i and σ_j instead of just one.²⁹ Also we stick to the step size of 0.5. Simple OLS regressions (see tables 18 and 19) show no significant time trend in any of the offers. Therefore

²⁹As before the equilibrium algorithm uses a fixed point iteration approach and consider a sequence to have converged when the sum of absolute probability difference of both vectors between the last two iterations is smaller than 0.001.

		Emp.	QRE
Offers	Principal A	1.20	1.81
	Principal B	0.90	0.92
Outcomes	UL	0.59	0.77
	UR	0.02	0.07
	DL	0.02	0.07
	DR	0.37	0.08

Table 10: Empirical and Best-Fit QRE Mean Total Offers and Outcomes in BoS

now we consider all periods 1 to 28.³⁰

As above we can now check if the equilibria are consistent with aggregate data patterns. Figures 11 and 12 again show histograms for empirically observed and best-fit QRE offers. Note that now that we consider the offers separately and thus ignore symmetry, individual offer patterns are very similar between observed and predicted offers. Tables 10 and 11 again juxtapose empirical and predicted mean offers as well as the distribution of outcomes. The predicted mean offers are quite close to the empirical means in both treatments. More importantly, principal A is predicted to offer more than principal B in either treatment, which is one feature of the empirical data. Finally, as mentioned before, in BoS the QRE cannot perfectly replicate the empirical distribution of offers because of the non-symmetric offers in the QRE. Still, the QRE predicts outcome UL to prevail most of the time, which we also observe in the data.

³⁰Results do not change much if we consider periods 15 to 28 as before.

		Emp.	QRE
Offers	U	1.12	1.31
	L	1.13	1.31
	D	0.70	0.68
	R	0.65	0.68
Outcomes	UL	0.72	0.80
	UR	0.10	0.09
	DL	0.08	0.09
	DR	0.09	0.02

Table 11: Empirical and Best-Fit QRE Mean Total Offers and Outcomes in DOM




Figure 11: Histograms of Empirical and Best-Fit QRE Offers in DOM




Figure 12: Histograms of Empirical and Best-Fit QRE Offers in BoS

Appendix C Additional Calculations, Tables and Figures

C.1 Equilibrium calculation DOM

Let us formally derive the equilibrium for this game. The unique efficient outcome in this game is UL, hence it must result in any pure-strategy Nash equilibrium. By *IC* we know that

$$5 \geq t_R^A(U) + t_C^A(L)$$

$$0 \geq 1 - t_R^B(D)$$

$$0 \geq 1 - t_C^B(R)$$

Obviously, the offers for D and R have to be at least 1 each and thus by **(AI)** the offers for U and L have to be at least 1 each but may not sum up to more than 5. The set of pure-strategy equilibria with UL as an outcome is

$$(t_R^A(U) = t_R^B(D), t_C^A(L) = t_C^B(R)) \in \{(a, b) \mid a + b \leq 5 \wedge a, b \geq 1\}$$

There exists a unique equilibrium with the lowest total transfer from the principals:

$$t_R^A(U) = t_C^A(L) = t_R^B(D) = t_C^B(R) = 1$$

C.2 Equilibrium calculation BoS

Let us now turn to treatment BoS. The unique efficient outcome in this game is UL, hence it must result in any pure-strategy Nash equilibrium. By *IC* we know that

$$2 + t_R^A(U) + t_C^A(L) \geq \max\{t_R^A(U), t_C^A(L), 4\}$$

$$5 \geq \max\{t_R^B(D), t_C^B(R), 2 + t_R^B(D) + t_C^B(R)\}$$

By the **(AI)** condition, in equilibrium, $t_R^A(U) = t_R^B(D)$ and $t_C^A(L) = t_C^B(R)$. Thus, the conditions above simplify to

$$5 \geq 2 + t_R^B(D) + t_C^B(L)$$

$$4 \leq 2 + t_R^A(U) + t_C^A(R)$$

and the set of pure-strategy equilibria with UL as an outcome is

$$(t_R^A(U) = t_R^B(D), t_C^A(L) = t_C^B(R)) \in \{(a, b) \mid a + b \in [2, 3]\}$$

There exists a unique *set* of equilibria with the lowest total transfer from the principals:

$$t_R^A(U) + t_C^A(L) = t_R^B(D) + t_C^B(R) = 2$$

Treatment	Half	Percent of Outcomes				Median Offer	
		UL	UR	DL	DR	U/L	D/R
PD High	1 st	0.52	0.24	0.20	0.04	1.90	1.09
PD High	2 nd	0.50	0.21	0.25	0.05	1.70	1.02
PD Low	1 st	0.59	0.18	0.18	0.04	1.26	0.50
PD Low	2 nd	0.63	0.17	0.14	0.05	1.00	0.40
COORD	1 st	0.88	0.06	0.04	0.01	0.00	0.00
COORD	2 nd	0.96	0.01	0.03	0.00	0.00	0.00

Table 12: Symmetric Treatments First and Second Half

Half	Percent of Outcomes				Median Offer			
	UL	UR	DL	DR	U	L	D	R
1 st	0.71	0.10	0.09	0.09	1.11	1.14	0.61	0.75
2 nd	0.60	0.02	0.02	0.36	1.03	1.05	0.69	0.60

Table 13: DOM First and Second Half

Half	Percent of Outcomes				Median Total Offer	
	UL	UR	DL	DR	Principal A	Principal B
1 st	0.57	0.02	0.03	0.38	1.00	0.75
2 nd	0.60	0.02	0.02	0.36	1.20	1.00

Table 14: BoS First and Second Half

Change	-	NA	+	Σ
	+	72	5	103
0	216	14	75	305
-	180	0	7	187
Σ	468	19	185	672

LDT
U/L PD High

Change	-	NA	+	Σ
	+	28	91	52
0	77	174	117	368
-	75	46	12	133
Σ	180	311	181	672

LDT
D/R PD High

Change	-	NA	+	Σ
	+	74	7	69
0	273	11	58	342
-	169	2	9	180
Σ	516	20	136	672

LDT
U/L PD Low

Change	-	NA	+	Σ
	+	23	58	28
0	68	293	112	473
-	41	44	5	90
Σ	132	395	145	672

LDT
D/R PD Low

Change	-	NA	+	Σ
	+	3	1	1
0	94	544	16	654
-	13	0	0	13
Σ	110	545	17	672

LDT
U/L COORD

Change	-	NA	+	Σ
	+	2	1	0
0	5	637	17	659
-	10	0	0	10
Σ	17	638	17	672

LDT
D/R COORD

Note: On the x-axis LDT shows learning direction theory's predictions. The theory either prescribes a lower offer (-), a higher offer (+) or doesn't yield a prediction (NA). On the y-axis Change shows whether a subject increased (+), decreased (-) or didn't change her offer (0) in the following period. Each cell counts the number of cases. Cells are colored green (dark gray when printed in black and white) if the theory correctly predicted the change, yellow (light gray when printed in black and white) if not, and white if it didn't yield a prediction. Here, we count zeros as successes.

Table 15: Learning Direction Theory Results in Periods 14-28

Offer for	PD High		PD Low		COORD	
	U/L	D/R	U/L	D/R	U/L	D/R
Period	-0.025** (0.010)	-0.018 (0.011)	-0.038*** (0.010)	-0.010 (0.008)	-0.043*** (0.007)	-0.009*** (0.003)
Constant	1.920*** (0.117)	1.246*** (0.133)	1.649*** (0.138)	0.715*** (0.107)	0.505*** (0.090)	0.129*** (0.046)
N	672	672	672	672	672	672
R^2	0.015	0.006	0.032	0.003	0.156	0.026

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Standard errors clustered on subject level in parentheses

Table 16: OLS Regressions With Periods 1-14

Offer for	PD High		PD Low		COORD	
	U/L	D/R	U/L	D/R	U/L	D/R
Period	0.000 (0.010)	-0.006 (0.011)	-0.005 (0.006)	0.005 (0.004)	-0.001 (0.001)	-0.002 (0.002)
Constant	1.551*** (0.235)	1.169*** (0.260)	1.155*** (0.159)	0.383*** (0.132)	0.027 (0.021)	0.050 (0.042)
N	672	672	672	672	672	672
R^2	0.000	0.001	0.001	0.001	0.011	0.009

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Standard errors clustered on subject level in parentheses

Table 17: OLS Regressions With Periods 15-28

Offer for	Principal A		Principal B	
	U	L	D	R
Period	-0.004 (0.005)	-0.001 (0.003)	0.001 (0.002)	0.001 (0.002)
Constant	0.667*** (0.123)	0.602*** (0.100)	0.435*** (0.059)	0.440*** (0.060)
N	1008	1008	1008	1008
R^2	0.003	0.000	0.000	0.000

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Standard errors clustered on subject level in parentheses

Table 18: OLS Regressions for BoS With Periods 1-28

Offer for	Principal A		Principal B	
	U	L	D	R
Period	0.003 (0.004)	-0.003 (0.003)	0.003 (0.004)	-0.003 (0.004)
Constant	1.079*** (0.082)	1.181*** (0.088)	0.666*** (0.091)	0.695*** (0.097)
N	672	672	672	672
R^2	0.002	0.004	0.001	0.003

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Standard errors clustered on subject level in parentheses

Table 19: OLS Regressions for DOM With Periods 1-28

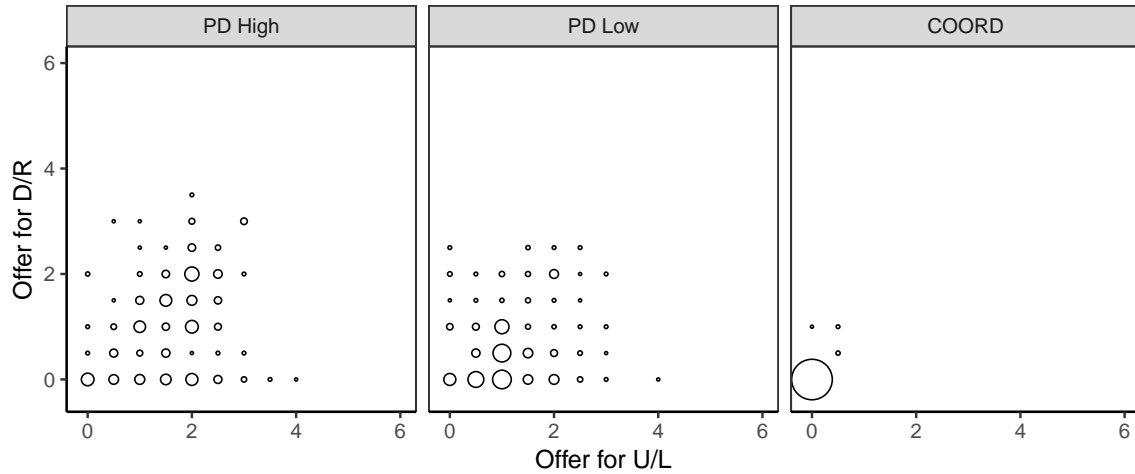



Figure 13: A Bubble Plot of Transfer Offers in Second Half



Note: The area of a bubble corresponds to the probability of the corresponding combination of offers. Very small bubbles may not be displayed.

Figure 14: QREs of Symmetric Treatments for Increasing λ

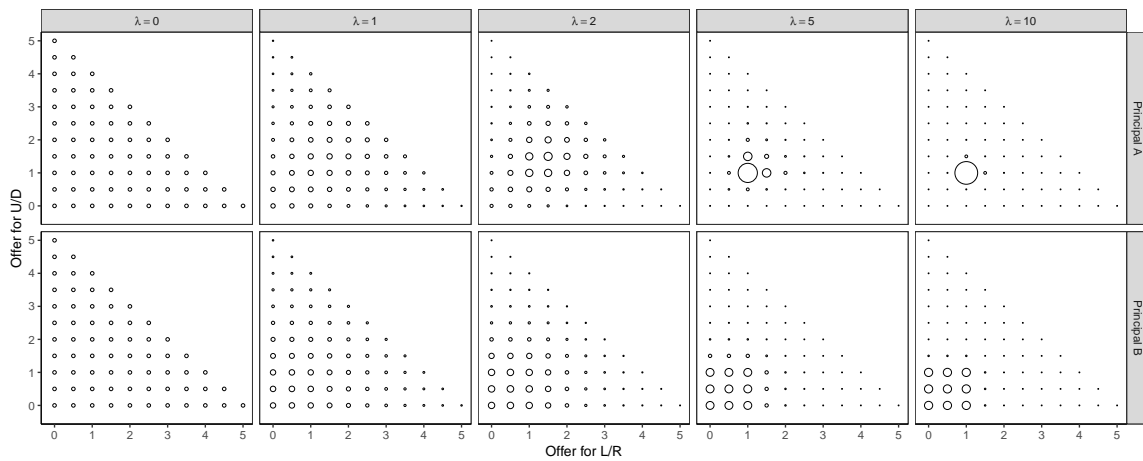



Figure 15: QREs of DOM Offers for Increasing λ



Note: The area of a bubble corresponds to the probability of the corresponding combination of offers. Very small bubbles may not be displayed.

Figure 16: QREs of BoS for Increasing λ

Appendix D Instructions

These instructions are translated from the German instructions that were used in the experiment.

Welcome to our experiment!

During the experiment you may not use electronic devices or communicate with other participants. Please use only the programs and functions that are provided for the experiment. Please do not talk to other participants. If you have a question, please raise your hand. We will come to you and answer your question in private. Please do not ask your questions out loud. If the question is relevant for all participants we will repeat and answer it aloud. If you break these rules we will have to exclude you from the experiment without pay.

Rules

In this experiment you and your opponent, who is another randomly chosen participant, will decide how to pay two computer agents whose actions determine your earnings in the experiment. The computer agents both have to decide between two actions. The first computer agent has to decide between Up and Down, the second one between Left and Right. The decisions made by the computer agents determine the earnings that you and your opponent will receive. The following table shows the four possible combinations of actions by the computer agents and your winnings (the exact procedure is described in section payment). L and R stand for Left and Right, U and D for Up and Down.

	L	R
U	4;4	0;6
D	6;0	1;1

You will receive the table's first component. Your opponent will receive the second component. Therefore,

- If the first computer agent chooses U and the second L, both you and your opponent will earn 4 Euro.
- If the first computer agent chooses U and the second R, you will earn 0 Euro and your opponent will earn 6 Euro.
- If the first computer agent chooses D and the second L, you will earn 6 Euro and your opponent will earn 0 Euro.
- If the first computer agent chooses D and the second R, both you and your opponent will earn 1 Euro.

Your opponent has the exact same information.

Computer Agents

Now you will wonder how the computer agents choose their action. They choose as follows: Both you and your opponent each can promise some amount of money to the computer agents in order to choose a specific action. For this both you and your opponent have a budget of 6 Euro.

You can promise the first computer agent some amount of money to choose Down. Your opponent can promise the first computer agent some amount of money to choose Up. How does the computer agent decide? She chooses the action for which she was promised more money. If she was promised the same amount for both actions, she will choose Up.

You can promise the second computer agent some amount of money to choose Left. Your opponent can promise the second computer agent some amount of money to choose Right. How does the computer agent decide? She too chooses the action

for which she was promised more money. If she was promised the same amount for both actions, she will choose Left.

[Note: The following paragraph was not included in treatment COORD since it did not apply there.]

Why do you offer money for D and L and your opponent for U and R? There is a reason for this. Note that it is always better for you if the first computer agent chooses D, no matter what the second computer agent does. (You will earn 6 Euro instead of 4 Euro if the second computer agent chooses L, respectively 1 Euro instead of 0 Euro if the second computer agent chooses R.) Similarly it is always better for you if the second computer agent chooses L. (You will earn 4 Euro instead of 0 Euro, respectively 6 Euro instead of 1 Euro.) In turn it is better for your opponent if the first computer agent chooses U and the second computer agent chooses R. This is why your opponent will promise money for these actions.

If a computer agent chooses an action that you promised money for, you will have to pay the amount, i.e. the promised amount is subtracted from your budget. If a computer agent chooses another auction, your opponent has to pay the amount of money she promised.

You can decide freely how much money to promise to the computer agents. The only condition is that the sum of the promised payments to both computer agents may not exceed your budget of 6 Euro. The same condition applies to your opponent.

Experimental Procedures

The situation that was described above will be repeated 30 times. Each repetition is one round. The first two rounds are trial rounds, while rounds 3 through 30 are relevant for your payouts at the end of the experiment (see section payment).

At the beginning of each round another participant will be randomly chosen by the computer and matched to you as your opponent.

At the beginning of each round you and your opponent will enter the payments that you promise to the computer agents for an action. After you and your opponent have entered the amounts the computer agents will choose their actions following the rules above. At the end of the round you will learn what payments your opponent promised and your net earnings from the round (see section payment).

Payment

Your net earnings in a round are calculated as follows. You start the round with a budget of 6 Euro. To this budget we will add the earnings that you get from the actions that the two computer agents choose (4 Euro for (U ; L), 0 Euro for (U ; R), 6 Euro for (D ; L), 1 Euro for (D ; R)). We will subtract the payments that you promised to the computer agents if they chose the corresponding action.

After the 30th and last round the computer will choose two out of the rounds 3 through 30 randomly with equal probability. Your net earnings from these rounds – and only the net earnings from these rounds – will be paid out to you in cash additionally to your participation fee.

Participation fee plus net earnings from these two rounds randomly chosen by the computer add up to the payments from these experiments, i.e. the amount of money that you will receive in cash from the experimental manager.

The money that the computer agents receive will not be paid out.

An Example

Assume that you offer 2 Euro to the first computer agent to choose D and 3 Euro to the second computer agent to choose L in the first round. Assume further that your opponent offers 2.50 Euro to the first computer agent to choose D and 2.70 to the second computer agent to choose R.

Then the first computer agent will choose U and the second computer agent will choose L. The outcome therefore is (U ; L). Therefore your net earnings in this round is

your budget of 6 Euro plus earnings of 4 Euro minus the payment that you promised to the second computer agent (because she choose L). This adds up to $6 + 4 - 3 = 7$ Euro. The net earnings of your opponent in this round therefore are $6 + 4 - 2.50 = 7.50$ Euro.

The numbers in this example have been chosen arbitrarily. They are only intended to illustrate the rules and procedures of the experiment. They are not a suggestion to you how to decide.

Comprehension Test

Please answer all 4 questions. You may continue with the experiment only after answering all questions correctly.

Assume that you offer 0.84 to the first computer agent to choose D and 3.50 to the second computer agent to choose L. Assume further that your opponent offers 1.73 Euro to the first computer agent to play U and 2.50 to the second computer agent to choose R.

1. What will the computer agents do?
2. What are your net earnings in this round?
3. What are your opponent's net earnings?
4. What would the second computer agent have done if you had offered her 2.50 Euro as well?

(The numbers in this example have been chosen arbitrarily. They are not a suggestion to you how to decide.)