DISCUSSION PAPER SERIES

No. 9510

DO SELLERS OFFER MENUS OF CONTRACTS TO SEPARATE BUYER TYPES? AN EXPERIMENTAL TEST OF ADVERSE SELECTION THEORY

Eva I Hoppe and Patrick W Schmitz

INDUSTRIAL ORGANIZATION

Centre for Economic Policy Research

www.cepr.org

Available online at: www.cepr.org/pubs/dps/DP9510.asp

ISSN 0265-8003

DO SELLERS OFFER MENUS OF CONTRACTS TO SEPARATE BUYER TYPES? AN EXPERIMENTAL TEST OF ADVERSE SELECTION THEORY

Eva I Hoppe, University of Cologne
Patrick W Schmitz, University of Cologne and CEPR

Discussion Paper No. 9510
June 2013

Centre for Economic Policy Research
77 Bastwick Street, London EC1V 3PZ, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

CEPR Discussion Paper No. 9510

June 2013

ABSTRACT

Do Sellers Offer Menus of Contracts to Separate Buyer Types? An Experimental Test of Adverse Selection Theory*

In the basic adverse selection model, a seller makes a contract offer to a privately informed buyer. A fundamental hypothesis of incentive theory is that the seller may want to offer a menu of contracts to separate the buyer types. In the good state of nature, total surplus is not different from the symmetric information benchmark, while in the bad state, private information may be welfare-reducing. We have conducted a laboratory experiment with 954 participants to test these hypotheses. While the results largely corroborate the theoretical predictions, we also find that private information may be welfare-enhancing in the good state.

JEL Classification: C72, C92, D82 and D86
Keywords: incentive theory, laboratory experiment, mechanism design and private information

Eva I Hoppe
University of Cologne
Albertus-Magnus-Platz
50923 Köln
GERMANY

Email: eva.hoppe@uni-koeln.de

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=169078

Patrick W Schmitz
University of Cologne
Albertus-Magnus-Platz
50923 Köln
GERMANY

Email: patrick.schmitz@uni-koeln.de

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=149900

*We would like to thank Felix Meickmann for providing excellent research assistance in programming and conducting the experiment.

Submitted 07 June 2013

1 Introduction

In the past three decades, the theory of contracts and incentives has been one of the most active fields of research in microeconomics.[1] In this paper, we report about a large-scale laboratory experiment designed to test basic hypotheses of incentive theory. Specifically, we examine subjects' behavior in a setting where a seller makes a contract offer to a buyer who has private information about his willingness-to-pay. This setting is often simply referred to as the standard "adverse selection" or "screening" problem and might be called the centerpiece of mechanism design theory.[2] A fundamental hypothesis of incentive theory is that the seller may want to induce separation between high-valuation and low-valuation buyers by offering a suitable *menu* of contracts. In the second-best solution, high-valuation buyers then consume the same quantity or quality as in a setting with symmetric information, while there is a downward distortion in the case of low-valuation buyers. Our aim is to study whether these predictions are supported by the data.

We consider the simplest possible adverse selection problem in which incentive theory predicts that the seller may want to separate buyer types by offering a menu of contracts. Specifically, suppose a seller can sell either good A or good B to a buyer. In line with the traditional mechanism design approach, the seller has full commitment power and makes a take-it-or-leave-it offer to the buyer. For simplicity, assume the seller has no costs. The buyer is either a low type or a high type with equal probability. A low-type buyer's valuation for a good is smaller than a high-type buyer's valuation for the good. Moreover, regardless of his type, a buyer's valuation for good A is always larger than his valuation for good B. Good B may thus be interpreted as a smaller

[1] For comprehensive textbook expositions of contract theory, see Laffont and Martimort (2002), Bolton and Dewatripont (2005), and Salanié (2005).

[2] While "adverse selection" originally referred to a potential consequence of asymmetric information, by now the term is usually used whenever a party has private information at the time the contract is written (while post-contractual information asymmetries, e.g. due to hidden actions, go under the heading of "moral hazard," see Maskin and Riley, 1984, and Hart and Holmström, 1987). Adverse selection models are at the heart of mechanism design theory (cf. Salanié, 2005). See Nobel Prize Committee (2007) for an appraisal of mechanism design in celebration of the pioneering contributions by Hurwicz, Maskin, and Myerson.

quantity or quality of the same product.

Under standard assumptions of rationality and profit-maximizing behavior, when there is symmetric information, then the parties will trade the efficient good A and the seller will extract the total gains from trade. However, suppose now that the buyer has private information about his type. Depending on the parameter constellation, it can be optimal for the seller to offer a menu of contracts such that a high-type buyer will purchase good A, while a low-type buyer will purchase good B. In particular, if the low-type buyer's valuation for good A is very small, the seller would have to set a very small price if she wanted to ensure trade of the efficient good A regardless of the buyer's type. It can then be more profitable for the seller to trade only the inefficient good B with a low-type buyer, which allows the seller to obtain a higher price for good A from a high-type buyer. Hence, the same good as under symmetric information is sold in the good state of nature (i.e., there is "no distortion at the top"), while in the bad state of nature only good B is sold (i.e., there is a downward distortion of the quantity or quality traded).

In our experimental study, we have conducted two private information treatments. In parameter constellation I, incentive theory predicts that the seller offers a menu of contracts to separate the buyer types. In parameter constellation II, the low-type buyer's valuation for good A is sufficiently large such that according to theory, the seller wants to trade good A with both buyer types. In addition, we have conducted two benchmark treatments which are similar to the two main treatments, except that there is symmetric information.

Results. Consider first parameter constellation I. It turns out that when the buyers have private information, the vast majority of sellers indeed offer an incentive-compatible menu of contracts. As a result, high-valuation buyers typically buy good A, while buying good B is the most frequent decision of low-valuation buyers. Comparing the private information treatment to the benchmark treatment with symmetric information, we find that total surplus levels do not differ in the good state of nature, while the total surplus levels are smaller under private information in the bad state of nature. Hence, the presence of private information is welfare-reducing. These results are all in good accord with the main hypotheses of adverse selection theory.

However, there are deviations. In particular, we observe that some offers are rejected and that the prices are on average smaller than predicted. These deviations occur already in the symmetric information benchmark treatments and are reminiscent of similar findings in the literature on ultimatum game experiments (see Güth et al., 1982).[3]

Next, consider parameter constellation II. Most sellers offer only good A, and also when a menu is offered buying good A is the most frequent decision of the buyers, regardless of their type. As predicted, in the bad state of nature, the total surplus levels do not differ between the private information treatment and the corresponding benchmark treatment with symmetric information. However, in the good state of nature, the total surplus levels are *larger* under private information than under symmetric information. In contrast to standard theory, the presence of private information can thus be welfare-improving.

A closer look at the data reveals that the latter finding is due to the fact that standard theory is too optimistic about the efficiency attained under symmetric information. Once we take into account that buyers tend to reject offers that would give them only a very small payoff, the welfare-enhancing effect of private information is actually a consequence of the fact that sellers' price-setting behavior is in line with adverse selection theory. In parameter constellation II, incentive theory predicts that under private information sellers set prices for good A that are small enough to make them acceptable for low-type buyers. In contrast, under symmetric information the sellers set larger prices when they know that the buyer has a high valuation. As a result, in the good state of nature there are less rejections when there is private information.

Across all four treatments, inspection of the data reveals that the vast majority of buyer decisions is compatible with standard preferences. However, other-regarding preferences might be useful to explain the observed deviations.

---

[3] Since Güth et al. (1982) conducted the first experiment on the ultimatum game some thirty years ago, it has become one of the most prominent games in experimental economics. In the ultimatum game, a proposer makes a take-it-or-leave-it offer regarding the division of a pie to a responder. If the responder rejects, both parties get zero. Very unequal divisions are often rejected, and on average proposers offer 30-40% of the pie. See Güth and Tietz (1990), Güth (1995), and Camerer (2003) for surveys.

We employ the quantal response equilibrium (QRE) approach developed by McKelvey and Palfrey (1995) to estimate structural models, taking into account that buyers may have other-regarding preferences. It turns out that other-regarding preferences are helpful to explain the data; however, they are on average less pronounced than is suggested in the literature on inequity aversion in ultimatum game experiments (e.g., Fehr and Schmidt, 1999). The fact that there is uncertainty about how a particular buyer of a given type will react to a specific offer makes the sellers' task more difficult. To better understand the sellers' behavior, we have conducted two additional control treatments in which the role of the buyer is played by the computer. The sellers know that the computer will respond to their offers as a profit-maximizer. In both parameter constellations it turns out that the fraction of sellers who offer a menu of contracts does not differ between the computer treatment and the original private information treatment. However, the sellers set larger prices in the computer treatments, coming much closer to the theoretically predicted offers. We thus conclude that in the original private information treatments a substantial amount of the deviations of the sellers' behavior from the theoretical predictions cannot simply be attributed to decision errors. The sellers understand that the buyers will not always react as profit-maximizers and adjust their offers accordingly.

Related literature. The theory of incentives is focused on the implications that the presence of private information has for the design of contracts. Empirical tests of incentive theory are thus impeded by an inherent data availability problem.[4] As has recently also been pointed out by Huck et al. (2011), controlled laboratory experiments can thus be particularly useful to directly test contract-theoretic models. So far, most experiments on contract theory (including Huck et al., 2011) have explored moral hazard problems.[5] To the best of our knowledge, the present study is the first direct experimental test of a

[4] See Prendergast (1999) and Chiappori and Salanié (2003) for reviews of empirical work on contract theory.

[5] This fact has recently also been emphasized by Cabrales et al. (2011), who experimentally investigate the effects of competition between privately informed agents. See also Asparouhova (2006) on experiments that study adverse selection in insurance markets with competition between lenders.

contract-theoretic model that captures the main features of the basic adverse selection problem as devised by Baron and Myerson (1982), Maskin and Riley (1984), and Guesnerie and Laffont (1984).[6]

In a previous study (Hoppe and Schmitz, 2013), we have conducted an experiment in a simpler setting in which a principal and an agent could either agree on the efficient trade level or not trade at all. In this setting, it was not possible to distinguish between a distortion that is predicted by theory and a rejection that is in contrast to theory.[7] Thus, the fundamental hypothesis of adverse selection theory that the principal may offer a menu of contracts to induce separation between agent types could not be tested in this simple setting. This is also the case in the earlier studies of pie-splitting games with private information.[8]

In a recent paper, Charness and Dufwenberg (2011) experimentally study a hidden information problem that was motivated by contract theory. However, these authors consider only the contract that according to standard theory would be optimal under symmetric information. They exogenously assume that the same contract is signed when there is hidden information and explore the effects of communication. Thus, their approach is orthogonal to contract theory, which studies how the design of contracts should be adapted when information asymmetries are introduced. In line with the contract-theoretic

[6]These papers build on the pioneering work by Mirrlees (1971); see also the closely related work by Mussa and Rosen (1978) and Goldman et al. (1984). The basic adverse selection problem has become a cornerstone of contract theory, see e.g. Fudenberg and Tirole (1991, ch. 7), Laffont and Martimort (2002, ch. 2), Bolton and Dewatripont (2005, ch. 2), and Salanié (2005, ch. 2). These books also survey numerous applications in areas such as monopolistic price-discrimination, public procurement, or regulation of natural monopolies (see also Laffont and Tirole, 1993).

[7]The focus of the experiment in Hoppe and Schmitz (2013) was on endogenous information acquisition in the spirit of Crémer and Khalil (1992), Lewis and Sappington (1997), and Kessler (1998).

[8]In the pie-splitting experiments studied by Forsythe et al. (1991), Kagel et al. (1996), and Harstad and Nagel (2004) the responder may have private information about the size of the pie. See also Mitzkewitz and Nagel (1993), Straub and Murnighan (1995), Croson (1996), Güth et al. (1996), Rapoport and Sundali (1996), Güth and van Damme (1998), and Huck (1999) for variants of the ultimatum game in which the proposer has private information about the size of the pie.

approach, in our experiment contracts are endogenously chosen and we study to what extent the predicted implications of private information on contracting are borne out by the data.

Organization of the paper. The remainder of the paper is organized as follows. In the next section, we introduce the theoretical framework that provides the starting point for our experimental study. We describe the experimental design in section 3 and we derive predictions in section 4. We present and analyze our experimental results in section 5. Finally, concluding remarks follow in section 6.

2 The theoretical framework

In this section, we develop the simplest conceivable model that captures the main features of the basic adverse selection problem. Consider a seller and a buyer, both of whom are risk-neutral. At an initial date 0, nature draws the buyer's type $\theta \in \{H, L\}$, where $prob\{\theta = H\} = \pi$.

The parties can trade a single indivisible unit of either good A or good B. The two goods can be interpreted as two different quantities or qualities of the same product. The buyer's valuation of good A is given by v_A^θ, while his valuation of good B is given by v_B^θ, where $0 < v_A^L < v_A^H$ and $0 < v_B^L < v_B^H$. Let $v_B^\theta < v_A^\theta$ for both types $\theta \in \{H, L\}$. Thus, good B represents a smaller quantity or lower quality of the product that can be traded. We assume that the Spence-Mirrlees single-crossing condition $v_A^H - v_B^H \geq v_A^L - v_B^L$ holds; i.e., the additional utility from consuming a larger quantity or better quality is larger for a high-type buyer than for a low-type buyer. Moreover, we assume that $v_B^L > \pi v_B^H$, which ensures that the seller will never completely exclude the low-type buyer from trade. For simplicity, let us suppose that the seller does not have any costs. The reservation utilities of both parties are zero. Note that in a first-best world, the parties would always agree to trade good A, regardless of the buyer's type.

At date 1, the seller makes a take-it-or-leave it offer to the buyer. She can offer only good A at price p_A, or only good B at price p_B, or she offers a menu (p_A, p_B). If the seller offers only one good, at date 2 the buyer can decide whether or not to buy the good at the stated price. If the seller offers a menu, then at date 2 the buyer can decide whether to buy good A or good B or no good.

Each party maximizes its (expected) payoff. If a party's profit-maximizing decision is not unique, then it chooses the one that is in the interest of the other party. All parameters of the model are common knowledge, with the possible exception of the buyer's type θ. Specifically, let us consider two different scenarios.

In the benchmark scenario, there is symmetric information; i.e., the seller knows the realization of the buyer's type θ when she offers a contract. In this case, the parties' equilibrium behavior is straightforward and summarized in the following proposition.

Proposition 1 *In the case of symmetric information, good A will be traded regardless of the buyer's type. The seller offers only good A and sets $p_A = v_A^\theta$ (equivalently, she may offer a menu with $p_A = v_A^\theta$ and $p_B \geq v_B^\theta$). The seller's profit is v_A^θ and the buyer's profit is 0.*

Hence, when there is symmetric information, the first-best outcome is achieved and the seller can extract the total surplus.

In the adverse selection scenario, the buyer has private information about his type; i.e., the seller does not know the realization of θ. In this case, equilibrium behavior depends on the parameter constellation.

Proposition 2 *Consider the case of private information.*

(i) Suppose $v_A^L < v_B^L + \pi(v_A^H - v_B^H)$. Then good A will be traded if $\theta = H$, while good B will be traded if $\theta = L$. The seller offers a menu with $p_A = v_A^H - (v_B^H - v_B^L)$ and $p_B = v_B^L$. The seller's expected profit is $\pi(v_A^H - (v_B^H - v_B^L)) + (1-\pi)v_B^L$ and the buyer's expected profit is $\pi(v_B^H - v_B^L)$.

(ii) Suppose $v_A^L \geq v_B^L + \pi(v_A^H - v_B^H)$. Then good A will be traded regardless of the buyer's type θ. The seller offers only good A and sets $p_A = v_A^L$ (equivalently, she may offer a menu with $p_A = v_A^L$ and $p_B \geq v_B^L$). The seller's profit is v_A^L and the buyer's expected profit is $\pi(v_A^H - v_A^L)$.

Proof. See the Appendix.

When the buyer has private information, the first-best solution will be achieved only if v_A^L is sufficiently large (case ii). To give the buyer an incentive to buy good A regardless of his type, the seller sets $p_A = v_A^L$, which implies that the high-valuation buyer gets a rent $v_A^H - v_A^L$.

Yet, if v_A^L is sufficiently small (case i), then the high-valuation buyer's rent would become very large, so the seller prefers to offer a menu which does not lead to the first-best outcome if the buyer has a low valuation. Specifically, the seller sets $p_B = v_B^L$, such that the low-valuation buyer will buy good B. To give the high-valuation buyer an incentive to buy good A, his rent when he buys good A must be at least $v_B^H - v_B^L$ (i.e., his rent if he bought good B). Hence, the seller sets $p_A = v_A^H - (v_B^H - v_B^L)$. Thus, there is a downward distortion in the quantity or quality traded if the buyer's valuation is low, while there is no distortion away from the first-best outcome if the buyer's valuation is high (i.e., there is "no distortion at the top").

3 Experimental design

Our experiment consists of six treatments. In four treatments, half of the participants in each session were randomly assigned to the role of sellers and the others to the role of buyers. Each of these treatments was run in six to seven sessions; each session had between 24 and 32 participants. In addition, there were two treatments in which the role of the buyers was played by the computer, so that all participants were in the role of sellers. Each of these two treatments was run in three sessions with 28 to 32 participants.

No subject was allowed to participate in more than one session. In total, 954 subjects participated in the experiment. All subjects were students of the University of Cologne from a wide variety of fields of study.[9] All interactions were anonymous; i.e., no subject knew the identity of its trading partner.

In order to ensure a large number of independent observations, each session consisted of only one round; i.e., there were no repetitions and this was known to the subjects. At the beginning of each session, written instructions were

[9]The computerized experiment was programmed and conducted with zTree (Fischbacher, 2007) and subjects were recruited using ORSEE (Greiner, 2004).

handed out to each subject.[10] We made use of the experimental currency unit ECU. At the end of each session, the players' payoffs were converted into euros (1 ECU = 0.12 €).[11]

Our two main treatments PI-I and PI-II are designed to explore whether the sellers will indeed offer menus of contracts to induce separation whenever they should do so according to adverse selection theory.

Private information treatment PI-I. Each seller is randomly matched with one buyer. There are two goods A and B. A buyer and a seller can trade at most one of these goods. Half of the buyers have high valuations ($v_A^H = 100$ and $v_B^H = 40$), while the other half of the buyers have low valuations ($v_A^L = 40$ and $v_B^L = 30$).

There are two stages. In the first stage, the seller decides whether to offer only good A, or only good B, or a menu with both goods. If the seller decides to offer only good A, she chooses a price p_A. If she decides to offer only good B, she chooses a price p_B. If she offers a menu with both goods, she chooses prices p_A and p_B. The prices can be any integer between 0 and 100. The seller does not know the buyer's type; all she knows is that it can be either high or low with equal probability ($\pi = 1/2$).

In the second stage, the buyer learns his type and then he makes his buying decision. If the seller has offered only one good, the buyer can either buy the good at the stated price or reject the offer. If the seller has offered a menu, the buyer can either buy good A at price p_A, or buy good B at price p_B, or reject the offer. The resulting profits are displayed in Table 1.

	Buyer's profit (high type)	Buyer's profit (low type)	Seller's profit
Buyer buys good A	$100 - p_A$	$40 - p_A$	p_A
Buyer buys good B	$40 - p_B$	$30 - p_B$	p_B
Buyer rejects	0	0	0

Table 1. The profits in parameter constellation I.

[10] The instructions for all treatments are in the Supplementary Material.

[11] In addition to the profits made in the experiment, subjects in the four treatments with human buyers were paid a participation fee of 70 ECU, which ensured total payoffs to be non-negative.

Private information treatment PI-II. This treatment is identical to the PI-I treatment, except that now $v_A^L = 70$. The profits are displayed in Table 2.

	Buyer's profit (high type)	Buyer's profit (low type)	Seller's profit
Buyer buys good A	$100 - p_A$	$70 - p_A$	p_A
Buyer buys good B	$40 - p_B$	$30 - p_B$	p_B
Buyer rejects	0	0	0

Table 2. The profits in parameter constellation II.

It is well-known from numerous experiments on the ultimatum game that responders tend to reject offers that would give a very large fraction of the pie to the proposer, and that in anticipation of such a rejection behavior proposers make relatively generous offers. For this reason, we conduct two benchmark treatments in which there is symmetric information. We can then compare our main treatments with the benchmark treatments in order to isolate the effects that the presence of private information has on contracting, which are the focus of the present paper.

Symmetric information treatment SI-I. This treatment is identical to the PI-I treatment, except that the seller knows the buyer's type when she makes her offer.

Symmetric information treatment SI-II. This treatment is identical to the PI-II treatment, except that the seller knows the buyer's type when she makes her offer.

Finally, the sellers' behavior in the main treatments may reflect their beliefs about the buyers' reactions. To see whether the sellers would offer the theoretically optimal contracts if they knew for sure that the buyers would respond as standard theory predicts, we have conducted two additional control treatments. In these treatments, the buyer's role is taken on by the computer.[12]

[12] On the use of computer treatments, see also the experimental studies by Houser and Kurzban (2002) and Huck et al. (2011).

Private information computer treatment PIC-I. This treatment constitutes a one-person decision problem. The first stage is identical to the first stage in the PI-I treatment, except that the seller knows that the role of the buyer will be played by the computer. The seller knows that in the second stage the computer will make his buying decision in order to maximize his profit (and that in case of indifference the computer will make the decision that is better for the seller).

Private information computer treatment PIC-II. This treatment is identical to the PIC-I treatment, except that now $v_A^L = 70$.

4 Predictions

Our primary interest is to explore how subjects behave in the *private information treatment PI-I.* According to Proposition 2(i), the seller induces separation of the buyer types by offering a menu with $p_A = 90$ and $p_B = 30$, such that good A is bought by a high-type buyer and good B is bought by a low-type buyer. Hence, the first-best solution is achieved if the buyer is of the high type (i.e., there is "no distortion at the top"), while the inefficient good B is traded if the buyer is of the low type (i.e., there is a "downward distortion" in the bad state of nature). Intuitively, as the low valuation for good A is relatively small, it is not profitable for the seller to set p_A so small that the buyer purchases good A regardless of his type. Thus, the seller is better off if she sells good B to the low-type buyer. The maximum price that a low-type buyer is willing to pay for good B is $p_B = 30$. Note that given this price, a high-type buyer would make a profit of 10 if he bought good B. As a consequence, the maximum price the seller will set for good A is $p_A = 90$, because otherwise a high-type buyer would buy good B.

Recall that the outcome predicted by standard contract theory assumes that it is common knowledge that all parties behave in a rational and profit-maximizing way. While in the light of previous experimental results we do not expect that the subjects' behavior will strictly adhere to these assumptions, we hypothesize that most sellers will indeed offer menus that induce most high-type buyers to buy the efficient good A and most low-type buyers to buy the inefficient good B.

According to our theoretical framework, the sellers will try to separate the buyers only if the low valuation for good A is relatively small. It is thus instructive to compare treatment PI-I with treatment PI-II. In the *private information treatment PI-II,* according to Proposition 2(ii), the seller offers only good A at the price $p_A = 70$ (equivalently, she could offer a menu with $p_A = 70$ and $p_B \geq 30$). Thus, the first-best outcome is always attained; i.e., the buyer buys good A regardless of his type. Intuitively, since in parameter constellation II the low valuation for good A is relatively large, it is profit-maximizing for the seller to set p_A small enough such that both buyer types are willing to buy good A.

While we expect deviations from the standard theory predictions, we hypothesize that indeed sellers will typically not try to separate buyers in the treatment PI-II, so that in most cases the efficient good A will be traded, regardless of whether the buyer's valuation is low or high.

The vast literature on ultimatum games has shown that subjects' behavior deviates from standard theory already when the size of the pie to be divided is commonly known. In order to properly assess whether the introduction of private information has the effects predicted by adverse selection theory, it is thus useful to compare the main treatments with the benchmark treatments that are identical to PI-I and PI-II except that there is symmetric information. In the *symmetric information treatment SI-I,* according to Proposition 1, the seller offers $p_A = 100$ if the buyer is of the high type, and she offers $p_A = 40$ if the buyer is of the low type. (Equivalently, she could offer menus with $p_A = 100$, $p_B \geq 40$ and $p_A = 40$, $p_B \geq 30$, respectively.) In the *symmetric information treatment SI-II,* according to Proposition 1, the seller offers $p_A = 100$ if the buyer is of the high type, and she offers $p_A = 70$ if the buyer is of the low type. (Equivalently, she could offer menus with $p_A = 100$, $p_B \geq 40$ and $p_A = 70$, $p_B \geq 30$, respectively.) Standard theory predicts that in the symmetric information treatments the first-best solution will always be attained; i.e., good A will be traded regardless of the buyer's type.

We hypothesize that indeed in the symmetric information treatments the efficient good A will be traded in most cases. In particular, we hypothesize that in PI-I the total surplus will be smaller than in SI-I in the case of low-type buyers (since private information implies a "downward distortion" in the bad state), while there will be no difference in the case of high-type buyers (since private information implies "no distortion at the top"). Finally, we hypothesize that the total surplus levels do not differ between PI-II and SI-II, regardless of the buyer's type.

5 Results

5.1 Overview

This section summarizes our central results. Tables 3 and 5 show descriptive statistics of our two main treatments in which the buyers have private information, while Tables 4 and 6 present descriptive statistics of the corresponding benchmark treatments with symmetric information.

	Menu	Only A	Only B		Total
# Obs. (Share)	74 (71.2%)	28 (26.9%)	2 (1.9%)		104
Mean p_A	58.41	54.79			57.41
Mean p_B	23.28		30.00		23.46
High type:					
Buy A	28/35 (80.0%)	14/16 (87.5%)			42/52 (80.8%)
Buy B	7/35 (20.0%)		1/1		8/52 (15.4%)
Reject	0/35 (0.0%)	2/16 (12.5%)	0/1		2/52 (3.9%)
Mean profit seller	48.74	45.81	30.00		47.48
Mean profit buyer	39.26	41.69	10.00		39.44
Mean total surplus	88.00	87.50	40.00		86.92
Low type:					
Buy A	3/39 (7.7%)**	3/12 (25.0%)**			6/52 (11.5%)**
Buy B	25/39 (64.1%)**		0/1		25/52 (48.1%)**
Reject	11/39 (28.2%)**	9/12 (75.0%)**	1/1		21/52 (40.4%)**
Mean profit seller	15.10**	8.25**	0.00		13.23**
Mean profit buyer	7.21**	1.75**	0.00		5.81**
Mean total surplus	22.31**	10.00**	0.00		19.04**

Table 3. Private information treatment PI-I. The stars indicate whether there are significant differences (* at the 5% level, ** at the 1% level) between the high and the low types.

Consider first Table 3, which shows the outcomes of the private information treatment PI-I. As predicted, the vast majority of sellers offered a menu. In these cases, high-valuation buyers typically bought good A, while almost two-thirds of the low-type buyers bought good B. As indicated in Table 3, the buying decisions differ significantly between the high- and low-valuation buyers.[13] Even if we take into consideration that about one out of four sellers offered only good A, buying good B was the most frequent decision among all low-valuation buyers.

At first sight, the total rejection rate of around 40% in the case of low-type buyers appears to be relatively large. However, note that if a menu is offered, then the rejection rate is only around 28%, which is not different from the rejection rate of low-type buyers when there is symmetric information (see Table 4). The fact that in PI-I there is a large rejection rate of low-type buyers when only good A is offered is not surprising, since given that only good A is offered, standard theory would predict an offer of $p_A = 100$, which would be accepted by high types only.

In line with our predictions, in PI-II and in the symmetric information treatments the vast majority of buyers bought good A, regardless of the buyer's type. Indeed, in each of these three treatments most sellers offered only good A.

Consider the symmetric information treatment SI-I (see Table 4). Regardless of whether the seller is matched with a low-type or a high-type buyer, the share of sellers who offer a menu is significantly smaller than in PI-I (p-values ≤ 0.001). Moreover, even when a menu was offered, buying good A was the most frequent decision, regardless of the buyer's type.

Similarly, in parameter constellation II (see Tables 5 and 6) the share of sellers who offer a menu is significantly smaller than in the PI-I treatment (p-values < 0.001). Even if a menu is offered, buying good A is always the most frequent decision.

[13] Throughout, we use two-tailed Mann-Whitney U tests in the case of prices, profits, and surplus levels, while we use two-tailed Fisher exact tests in the case of categorical data.

	Menu	Only A		Total
High type:				
# Obs. (Share)	19 (40.4%)	28 (59.6%)		47
Mean p_A	67.00	70.18		68.89
Mean p_B	26.53			26.53
Buy A	16/19 (84.2%)	23/28 (82.1%)		39/47 (83.0%)
Buy B	2/19 (10.5%)			2/47 (4.3%)
Reject	1/19 (5.3%)	5/28 (17.9%)		6/47 (12.8%)
Mean profit seller	56.47	54.00		55.00
Mean profit buyer	31.95	28.14		29.68
Mean total surplus	88.42	82.14		84.68
Low type:				
# Obs. (Share)	17 (36.2%)	30 (63.8%)		47
Mean p_A	28.18**	29.53**		29.04**
Mean p_B	25.29			25.29
Buy A	8/17 (47.1%)*	22/30 (73.3%)		30/47 (63.8%)
Buy B	4/17 (23.5%)			4/47 (8.5%)
Reject	5/17 (29.4%)	8/30 (26.7%)		13/47 (27.7%)
Mean profit seller	16.59**	18.70**		17.94**
Mean profit buyer	9.29**	10.63**		10.15**
Mean total surplus	25.88**	29.33**		28.09**

Table 4. Symmetric information treatment SI-I. The stars indicate whether there are significant differences (* at the 5% level, ** at the 1% level) between the high and the low types.

So far, the main outcomes of the experiment are in good accord with the qualitative predictions of adverse selection theory. Yet, there are deviations. In all four treatments, we observe rejections, which should not occur according to standard theory. Moreover, on average the prices are smaller than predicted. As a consequence, the sellers are typically worse off and the buyers are better off compared to the theoretical benchmark. These deviations are reminiscent of experimental findings in the literature on ultimatum games. Since our goal is to explore whether incentive theory captures the main effects that occur when private information is added to a contracting problem, it is crucial to compare our private information treatments with the corresponding symmetric information treatments (which closely resemble standard ultimatum games).

When we compare the PI-I treatment with the relevant benchmark SI-I, we find strong support for the "no distortion at the top" prediction of adverse selection theory: The total surplus levels in case of high-type buyers do not differ significantly between PI-I and SI-I (p-value=0.958). Furthermore, in line with the predicted distortion in the bad state of the world, in the case of low-type buyers the total surplus level in PI-I is significantly smaller than in SI-I (p-value$<$0.001). Overall, as predicted, the impact of private information in parameter constellation I is thus to reduce the expected total surplus.

	Menu	Only A	Only B	Total
# Obs. (Share)	34 (37.0%)	56 (60.9%)	2 (2.2%)	92
Mean p_A	59.74	54.89		56.72
Mean p_B	28.97		27.50	28.89
High type:				
Buy A	17/18 (94.4%)	26/28 (92.9%)		43/46 (93.5%)
Buy B	0/18 (0.0%)			0/46 (0.0%)
Reject	1/18 (5.6%)	2/28 (7.1%)		3/46 (6.5%)
Mean profit seller	53.11	48.89		50.54
Mean profit buyer	41.33	43.96		42.93
Mean total surplus	94.44	92.86		93.48
Low type:				
Buy A	10/16 (62.5%)*	22/28 (78.6%)		32/46 (69.6%)**
Buy B	1/16 (6.3%)		1/2	2/46 (4.4%)
Reject	5/16 (31.3%)	6/28 (21.4%)	1/2	12/46 (26.1%)*
Mean profit seller	35.94	42.61	12.50	38.98
Mean profit buyer	9.69**	12.39**	2.50	11.02**
Mean total surplus	45.63**	55.00**	15.00	50.00**

Table 5. Private information treatment PI-II. The stars indicate whether there are significant differences (* at the 5% level, ** at the 1% level) between the high and the low types.

Now consider parameter constellation II. According to adverse selection theory, the total surplus levels should not differ between the private information and the symmetric information scenarios. Indeed, there is no significant difference in the case of low-type buyers (p-value=0.517). Yet, we find that the surplus level in the case of high-type buyers in PI-II is significantly larger than in SI-II (p-value=0.041). Hence, in parameter constellation II, the presence of private information surprisingly improves the total surplus, which could never happen under standard theory.

To better understand the decisions made by the subjects that have led to these results, we will analyze the data in more detail in the following section.

	Menu	Only A		Total
High type:				
# Obs. (Share)	14 (29.8%)	33 (70.2%)		47
Mean p_A	69.64	67.76		68.32
Mean p_B	37.57			37.57
Buy A	12/14 (85.7%)	25/33 (75.8%)		37/47 (78.7%)
Buy B	0/14 (0.0%)			0/47 (0.0%)
Reject	2/14 (14.3%)	8/33 (24.2%)		10/47 (21.3%)
Mean profit seller	56.07	47.94		50.36
Mean profit buyer	29.64	27.82		28.36
Mean total surplus	85.71	75.76		78.72
Low type:				
# Obs. (Share)	9 (19.2%)	38 (80.9%)		47
Mean p_A	45.44**	53.11**		51.64**
Mean p_B	30.33			30.33
Buy A	6/9 (66.7%)	30/38 (78.9%)		36/47 (76.6%)
Buy B	0/9 (0.0%)			0/47 (0.0%)
Reject	3/9 (33.3%)	8/38 (21.1%)		11/47 (23.4%)
Mean profit seller	24.44**	39.32*		36.47**
Mean profit buyer	22.22	15.95**		17.15**
Mean total surplus	46.67**	55.26**		53.62**

Table 6. Symmetric information treatment SI-II. The stars indicate whether there are significant differences (* at the 5% level, ** at the 1% level) between the high and the low types.

5.2 A closer look at the data

Figure 1 shows the distribution of the sellers' offers in the private information treatments. In each treatment, the size of a circle is proportional to the relative frequency with which the respective offer was made.[14]

First, consider the treatment PI-I, which is displayed in the left panel of Figure 1. Recall that 74 sellers offered a menu. Their offers (p_A, p_B) are shown in the lower part of the PI-I panel. The offers p_A of the 28 sellers who offered good A only are shown in the upper part of the panel. The most frequently observed offer was a menu with $p_A = 70$, $p_B = 20$; this offer was made by 8 sellers.[15]

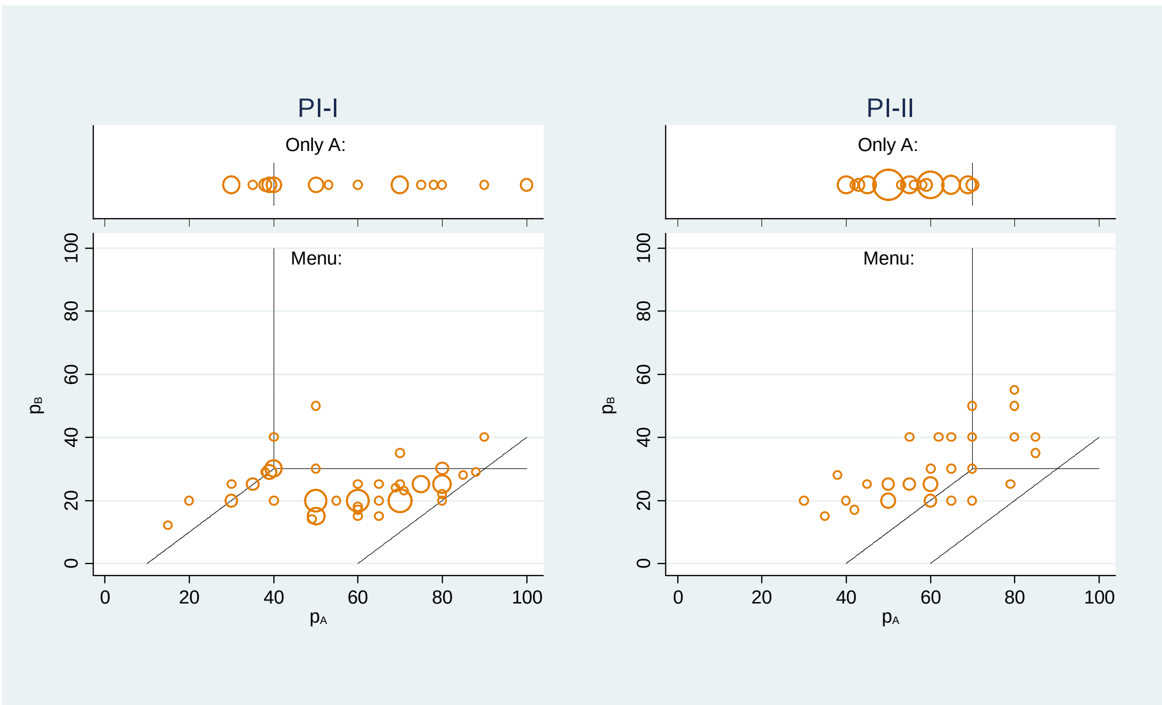

Figure 1. The distribution of the sellers' offers in the treatments PI-I and PI-II. In each treatment, the size of the circles is proportional to the relative frequency with which the offers were made.

In parameter constellation I, according to standard theory low-type buyers never buy good A if p_A is larger than 40, and they never buy good B if p_B is larger than 30. These critical prices are illustrated in the PI-I panel by

[14]In each of the two private information treatments, there were two sellers who offered good B only. These offers are not depicted in the figures.

[15]The other offers that were made by at least five sellers were menus with $p_A = 60$, $p_B = 20$ (7 times), $p_A = 50$, $p_B = 20$ (7 times), and $p_A = 80$, $p_B = 25$ (5 times).

the vertical and horizontal black lines, respectively. Moreover, when a menu is offered, a low-type buyer prefers good B to good A whenever the incentive compatibility constraint $30-p_B \geq 40-p_A$ is satisfied; i.e. if the offer lies on the right-hand side of the curve $p_B = p_A - 10$. For a high-type buyer, purchasing good A is always more attractive than rejecting an offer $p_A \leq 100$. When a menu is offered, a high-type buyer prefers good A to good B whenever the incentive compatibility constraint $100 - p_A \geq 40 - p_B$ is satisfied, which is the case for offers that lie on the left-hand side of the curve $p_B = p_A - 60$.

Observe that according to standard theory, menu offers with $p_B \leq 30$ that satisfy both incentive compatibility constraints induce high-type buyers to buy good A and low-type buyers to buy good B. As is illustrated in Figure 1, 66 of the 74 menu offers (89.2%) lie in this region. Thus, a fundamental prediction of adverse selection theory is corroborated by our data. The sellers offer menus in order to separate the buyers depending on their types.

Next, consider the treatment PI-II, which is displayed in the right panel of Figure 1. As can be seen in the figure, the most frequently observed decision of the sellers was to offer only good A at the price $p_A = 50$; this offer was made 14 times.[16] In parameter constellation II, according to standard theory low-type buyers do not buy good A if p_A is larger than 70, and they do not buy good B if p_B is larger than 30. When a menu is offered, low-type buyers prefer good A to good B whenever the incentive compatibility constraint $70 - p_A \geq 30 - p_B$ is satisfied (which is the case if the offer lies on the left-hand side of the curve $p_B = p_A - 40$). When a menu is offered to a high type, he prefers good A to good B if $100 - p_A \geq 40 - p_B$, which holds for offers that lie left of the curve $p_B = p_A - 60$. In total, 82 of the 92 offers (89.1%) were such that a buyer should have bought good A regardless of his type. Thus, a comparison between PI-I and PI-II shows that in line with adverse selection theory, the sellers try to separate the buyers only if the low-type buyers' valuation for good A is sufficiently small.

Let us now examine the buyers' responses to the offers made by the sellers. The colors of the circles in Figure 2 show the buyers' behavior in treatments PI-I and PI-II. The left (right) panels depict all offers received by low-type

[16] The other offers that were made by at least five sellers were $p_A = 60$ (11 times) and $p_A = 65$ (5 times).

(high-type) buyers and their responses. As explained above, the black lines divide the offer space into different regions, depending on whether according to standard theory a low-type (high-type) buyer would buy good A, good B, or reject the offer (R). For the moment, ignore the orange lines and letters.

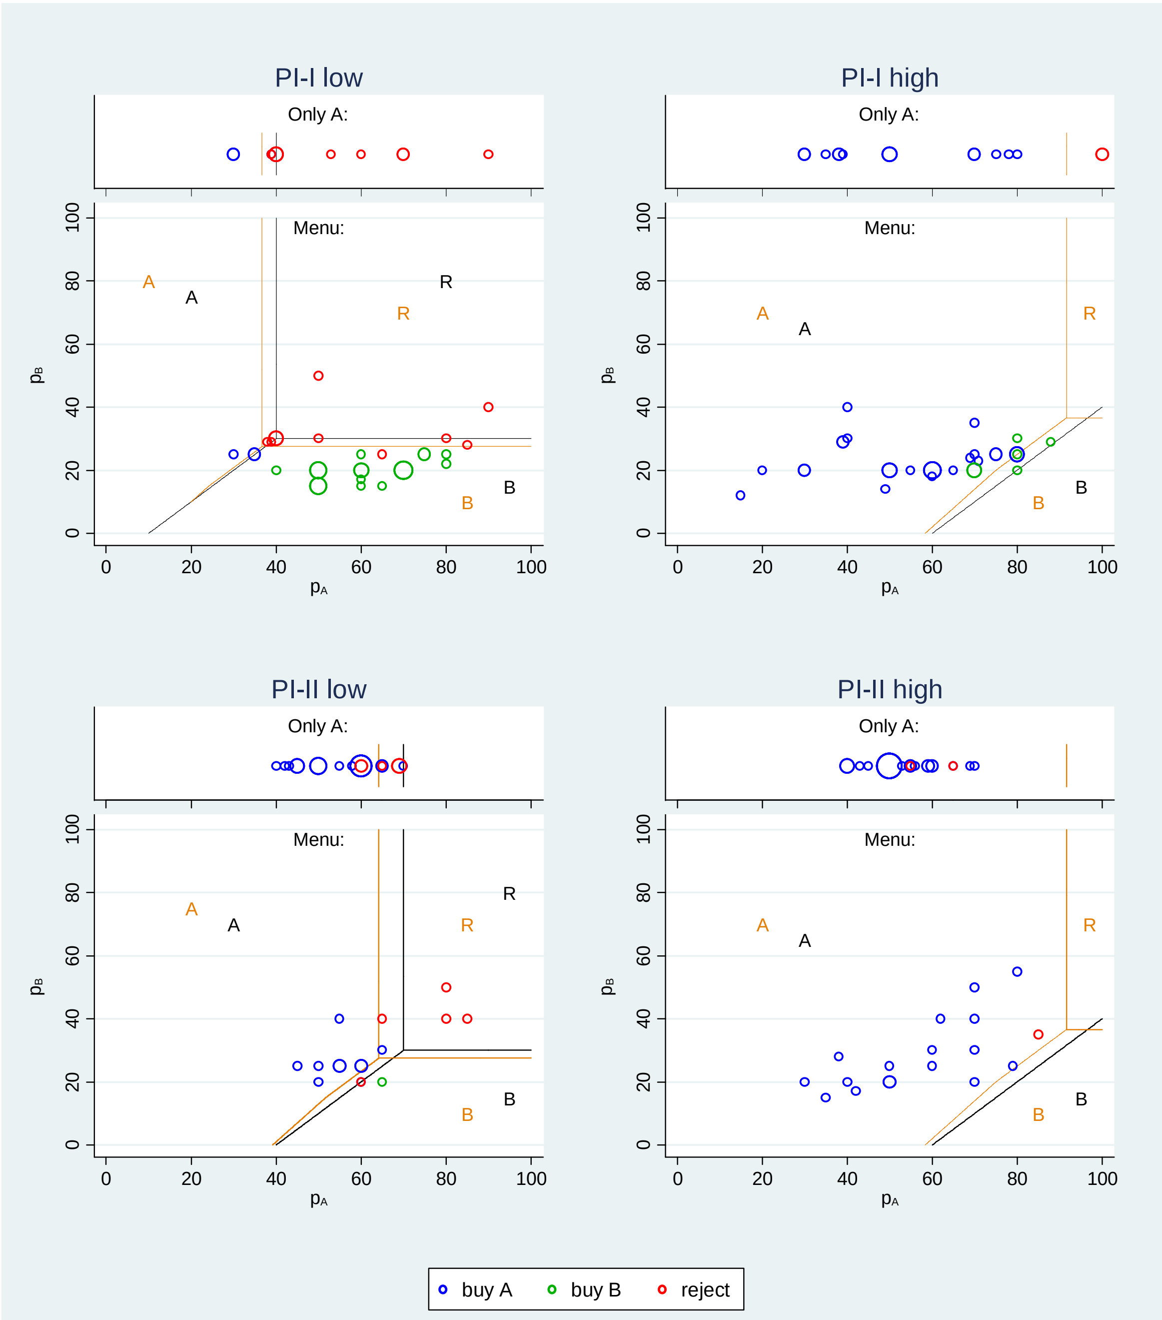

Figure 2. The buyers' reactions to the sellers' offers in the treatments PI-I and PI-II, depending on the buyers' type. The size of the circles is proportional to the relative frequency with which the offers were made.

The figure clearly shows that while there were some deviations (typically close to the borders between the regions), by far most buyer decisions were as expected; i.e., most circles in region "A" ("B") are indeed blue (green), while most circles in region "R" are red.

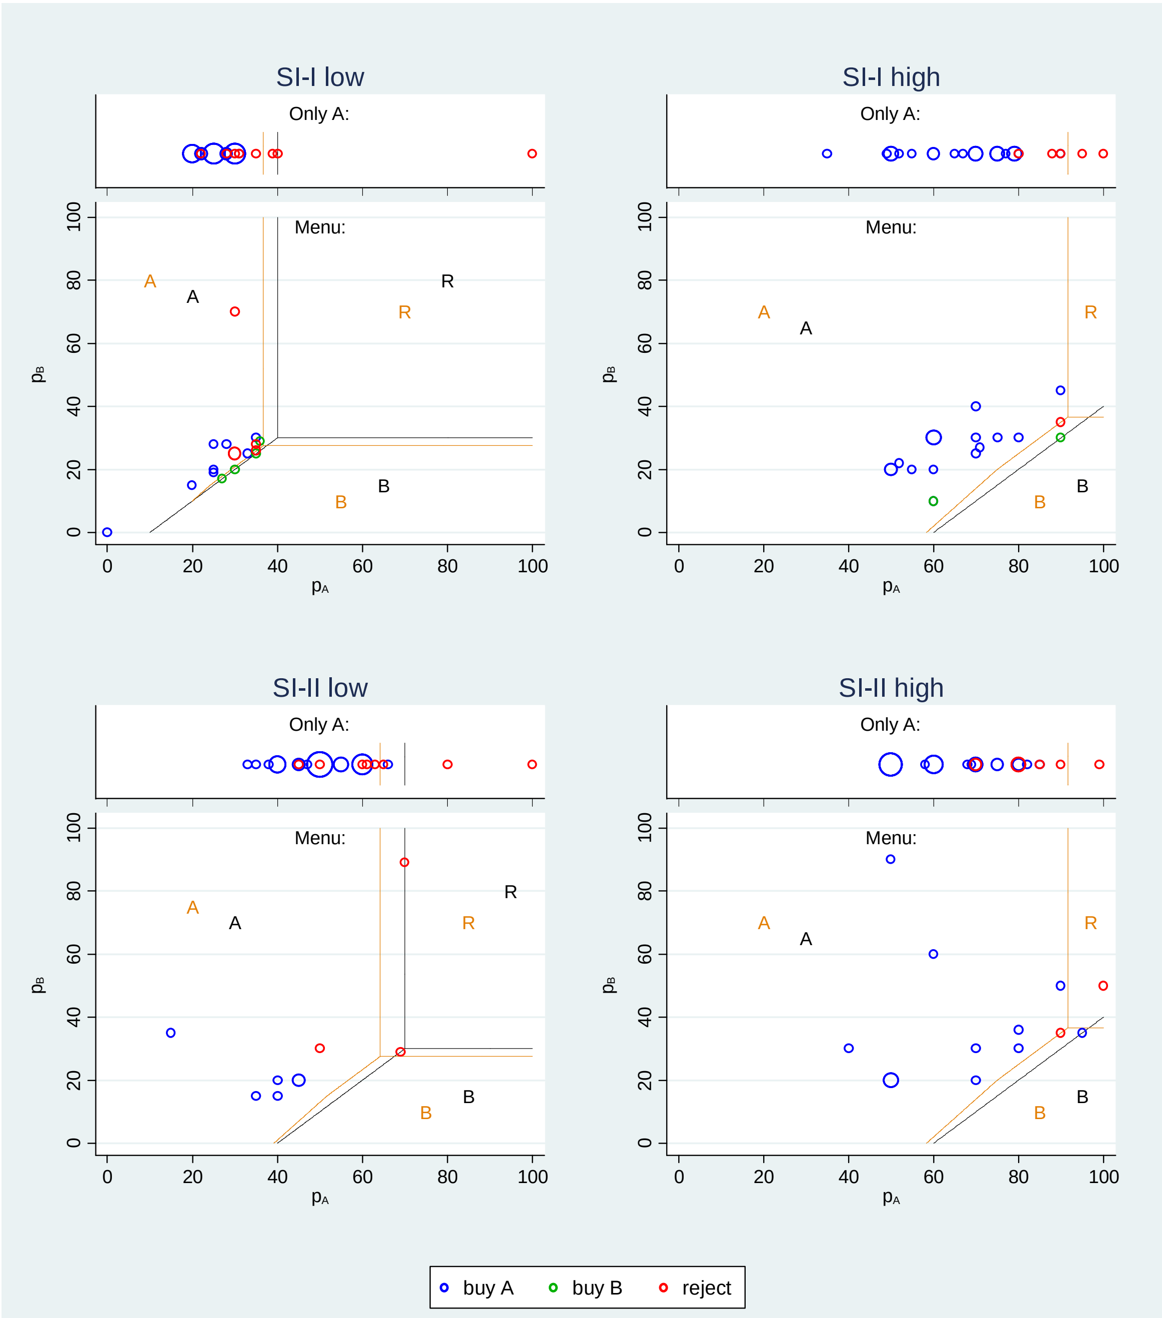

Figure 3. The sellers' offers and the buyers' reactions in the treatments SI-I and SI-II, depending on the buyers' type. The size of the circles is proportional to the relative frequency with which the offers were made.

Next, let us investigate the sellers' and buyers' behavior in the benchmark treatments with symmetric information. In these treatments, each seller knows whether she is matched with a high-type or a low-type buyer. Figure 3 shows the sellers' offers and the buyers' responses, depending on the buyers' type. Observe that regardless of the parameter constellation and regardless of the buyers' type, almost all sellers make offers such that according to standard theory buyers should buy good A. Indeed, the vast majority of the buyers responded as predicted, although there were some deviations.

When we compare the upper left panels of Figures 2 and 3, which depict the offers received by low-valuation buyers in PI-I and SI-I, it becomes evident that in PI-I much more offers induced the buyers to purchase the inefficient good B. As can be seen in the upper right panels of Figures 2 and 3, the vast majority of high-valuation buyers bought good A, regardless of whether or not there was private information. From an ex ante point of view, we find that in line with adverse selection theory private information reduces efficiency in parameter constellation I.

Now let us examine parameter constellation II. As can be seen in the lower left panels of Figures 2 and 3, behavior did not differ much between PI-II and SI-II when the buyers had low valuations. However, compare now the lower right panels, which show the offers received by high-type buyers in PI-II and SI-II. Observe that in PI-II, only 3 out of the 46 sellers (i.e., less than 7%) chose a price p_A strictly larger than 70. In line with adverse selection theory, in parameter constellation II the sellers wanted to trade good A with both types of buyers. Since in PI-II the seller did not know the buyer's type, she did not set p_A larger than 70 if she wanted to trade good A regardless of the buyer's type. In contrast, when in SI-II the sellers knew that the buyers' valuation for good A was 100, then 18 out of 47 sellers (i.e., more than 38%) chose a price p_A strictly larger than 70. Yet, several of these offers were rejected, whereas in PI-II there were fewer offers with $p_A > 70$ and thus there were fewer rejections. This explains why overall the introduction of private information enhances efficiency in parameter constellation II (whereas under standard theory, full efficiency would already be attained under symmetric information, so there would be no scope for an efficiency increase due to private information).

Taken together, standard theory is too optimistic about the efficiency achieved in the symmetric information benchmark, since it neglects the fact that offers which give buyers only a very small fraction of the surplus tend to be rejected. Once we take into account such rejection behavior, the conclusions of adverse selection theory about the sellers' behavior when buyers have private information are corroborated by the data. In particular, in parameter constellation I most sellers offer menus to induce separation, while in parameter constellation II the sellers want to trade good A with both buyer types.

5.3 Analyzing the deviations

In this section, we analyze the quantitative deviations of the data from the theoretical predictions more deeply. In the literature on ultimatum games, it has often been argued that rejections of offers that give responders only a small fraction of the pie can be explained by other-regarding preferences. In our framework, maybe the simplest formalization of such preferences is to assume that a buyer of type $\theta \in \{H, L\}$ has the utility $v_G^\theta - p_G - \alpha \max\{p_G - (v_G^\theta - p_G), 0\}$ if he buys good $G \in \{A, B\}$, and he has the utility zero if he buys no good. In other words, if in the case of a purchase the buyer's material payoff $v_G^\theta - p_G$ is larger than the seller's material payoff p_G, then the buyer's utility is equal to his material payoff. Otherwise, the buyer experiences a loss of α times the difference between the seller's and his own material payoff. The parameter $\alpha \geq 0$ measures the strength of the buyer's inequity aversion. The linear formalization is in the spirit of Fehr and Schmidt's (1999) prominent work on other-regarding preferences.[17]

It is straightforward to see that an inequity-averse buyer will prefer buying good G at price p_G to a rejection whenever $p_G \leq v_G^\theta(1+\alpha)/(1+2\alpha)$. Thus, while a buyer with standard preferences would be willing to buy good G whenever $p_G \leq v_G^\theta$, a sufficiently inequity-averse buyer may prefer to reject an offer $p_G \in (v_G^\theta/2, v_G^\theta]$. Analogously, an inequity averse-buyer may purchase the inefficient good B even when $v_A^\theta - p_A \geq v_B^\theta - p_B$ holds.

In order to estimate the parameter α, we must somehow account for noise

[17] See also the literature survey by Fehr and Schmidt (2006) for a discussion of related formalizations of other-regarding preferences.

in the data. One possibility to do so is the logit-QRE approach pioneered by McKelvey and Palfrey (1995). Specifically, let U_{ik} denote player i's expected utility if he makes a decision $k \in \{1, ..., n\}$. Then the probability that he chooses the decision $k = \hat{k}$ is given by

$$\frac{e^{\lambda U_{i\hat{k}}}}{\sum_{k=1}^{n} e^{\lambda U_{ik}}}.$$

Hence, the probability of a decision error decreases with the utility loss that is caused by the error. When computing his expected utility, a player takes into account that all decisions (including his own future decisions) are made in this stochastic way (we thus apply the agent-QRE concept developed by McKelvey and Palfrey, 1998). The parameter λ can be interpreted as a rationality parameter. Behavior is completely random if $\lambda = 0$, while behavior approaches rational choice when λ increases. We use maximum likelihood to estimate the model parameters. Following Rogers et al. (2009), we provide two benchmarks. The "random" log likelihood is a lower bound for the quality of fit; it results from a model where all decisions are randomly taken. The "empirical" log likelihood is the best possible fit to the aggregate data; it results from a hindsight model that assigns to each decision its empirical relative frequency.

Table 7 shows the results of the QRE estimations.[18] In the first row, we assume standard preferences, while in the second row, we allow each buyer to be inequity-averse with parameter α. Note that in the latter case it turns out that in every treatment the parameter α is significantly different from zero and the quality of fit is better than in the case of standard preferences. Yet, consistent among all four treatments we find that the quality of fit improves only modestly and the inequity aversion parameter is rather small, $\alpha \approx 0.1$.[19]

[18] In the QRE estimations, we have grouped the prices into 21 categories (specifically, category 0 contains prices weakly smaller than 2, categories $\tilde{p} \in \{5, 10, 15, ..., 95\}$ contain prices from $\tilde{p} - 2$ to $\tilde{p} + 2$, and category 100 contains prices weakly larger than 98).

[19] Based on their inspection of ultimatum game data, Fehr and Schmidt (1999) suggest an average α of 0.85 (without employing a formal estimation technique). Specifically, they propose a distribution of inequity aversion parameters also capturing aversion towards advantageous inequity. We have also estimated a QRE model assuming their distribution of parameters. In all four treatments, the quality of fit turns out to be lower than under the assumption of standard preferences.

	PI-I	PI-II	SI-I	SI-II
Standard preferences	$\lambda = 0.162$ (0.011) $\ln L = -611.4$	$\lambda = 0.128$ (0.011) $\ln L = -431.5$	$\lambda = 0.117$ (0.008) $\ln L = -466.7$	$\lambda = 0.090$ (0.006) $\ln L = -415.8$
Inequity aversion	$\alpha_A = 0.100$ (0.026) $\lambda = 0.162$ (0.011) $\ln L = -601.1$	$\alpha_A = 0.116$ (0.025) $\lambda = 0.135$ (0.011) $\ln L = -417.3$	$\alpha_A = 0.095$ (0.024) $\lambda = 0.122$ (0.008) $\ln L = -457.1$	$\alpha_A = 0.090$ (0.026) $\lambda = 0.093$ (0.006) $\ln L = -409.1$
Random	$\ln L = -758.3$	$\ln L = -562.2$	$\ln L = -578.8$	$\ln L = -534.0$
Empirical	$\ln L = -357.2$	$\ln L = -297.5$	$\ln L = -340.9$	$\ln L = -325.0$

Table 7. QRE estimations with standard preferences and with inequity aversion. (Standard errors are in parentheses.)

In the four treatments, altogether 384 buyer decisions had to be made. Only 57 decisions could not be explained by standard preferences; i.e., 85.2% of the decisions were compatible with standard theory. The regions in which good A, good B, or no good would be bought by an inequity-averse buyer with $\alpha = 0.1$ are delineated in Figures 2 and 3 by the orange curves (and they are labeled with the orange letters "A", "B", and "R"). For instance, in the case of low-type buyers in PI-I, some rejections that are incompatible with standard preferences can be explained if we allow for inequity aversion. Analogously, in the case of high-type buyers in PI-I, a few choices of good B that are incompatible with standard preferences can be explained by inequity aversion. Yet, there are also purchases of good A that could be explained by standard theory, while they are incompatible with inequity aversion (for example, see the case of low-type buyers in PI-II). Overall, 86.7% of the buyer decisions can be explained by inequity aversion with the parameter $\alpha = 0.1$, which is slightly better than standard theory.

The buyers' behavior typically deviated from standard predictions only relatively close to the borders between the regions in which good A, good B, or no good should be bought according to standard theory. This means that the utility loss associated with these deviations was rather small, so they can well be explained by QRE even in the absence of inequity aversion. Therefore, the QRE approach might somewhat underestimate the parameter α.

However, even without resorting to QRE estimations, there is evidence that other-regarding preferences are not very pronounced in our experiment.

In particular, we find that the fraction of buyer decisions that can be explained by inequity aversion is maximal for $\alpha = 0.2$; in this case, 87.5% of the decisions can be explained. For values of α larger than 0.38, less than 80% of the buyer decisions can be explained; for α larger than 0.67, less than 70% of the buyer decisions can be explained.

Of course, if we allowed each buyer to have a different parameter α, inequity aversion could explain more buyer decisions. In the QRE estimation, the parameter λ captures both noise and heterogeneity among players.[20] Overall, our analysis shows that while other-regarding preferences can be helpful to explain the buyers' behavior, on average these preferences are not as strong as suggested by the literature on inequity aversion in the context of ultimatum games.[21] Yet, buyers may have different inequity aversion parameters and they may make mistakes, so there is uncertainty about how a specific buyer of a given type will react to a particular offer, which makes the sellers' task more difficult.

In view of this analysis, it is not surprising that the sellers' offers are somewhat dispersed. Given that buyers may be inequity-averse, it makes sense for the sellers to offer smaller prices than predicted by standard theory. Given that a buyer's behavior is uncertain, different sellers may form different beliefs about what is the optimal price offer. However, it might also be the case that the sellers are simply not able to find the optimal contract, even if they knew for sure that buyers responded according to standard theory.

To shed more light on the sellers' behavior, we have conducted two control treatments which were similar to PI-I and PI-II, except that the role of the buyer was played by the computer. All participants in the treatments PIC-I and PIC-II knew that the computer would react to their offers as a profit

[20] See also De Bruyn and Bolton (2008) and Blanco et al. (2011), who also use the QRE approach to estimate other-regarding preferences.

[21] Most studies of ultimatum games use a divide-the-pie framing, making payoff comparisons particularly salient. Hoffman et al. (1994) have conducted treatments with a divide-the-pie framing as well as a seller-buyer framing, and they have found that results in the latter case are closer to standard theory. If we employ the same QRE approach as above to their data, we find $\alpha \approx 0.65$ in their divide-the-pie treatment, while $\alpha \approx 0.26$ in their seller-buyer treatment. Our setting with two goods may have further shifted the participants' attention to strategic considerations, thus making relative payoffs less salient.

maximizer. The descriptive statistics of the computer treatments are summarized in Tables 8 and 9. In both parameter constellations, the shares of sellers who offer a menu or only good A do not differ significantly between PI and PIC. Thus, as in PI-I, the vast majority of the participants in PIC-I offered a menu. Yet, the prices p_A and p_B were significantly larger than in PI-I (p-values<0.001). As in PI-II, most participants in PIC-II offered only good A, but the prices p_A were significantly larger than in PI-II (p-value<0.001).

The distributions of the offers that were made in the computer treatments are depicted by the black circles in Figure 4. For comparison, the orange circles show the offers in the original private information treatments. As explained above, the black (resp., orange) curves again delineate the regions in which low- and high valuation buyers buy good A, good B, or no good according to standard theory (resp., inequity aversion theory with $\alpha = 0.1$).

	Menu	Only A		Total
# Obs. (Share)	72 (75.0%)	24 (25.0%)		96
Mean p_A	79.22	72.92		77.65
Mean p_B	30.57			30.57
High type:				
Buy A	36/38 (94.7%)	10/10 (100.0%)		46/48 (95.8%)
Buy B	2/38 (5.3%)			2/48 (4.2%)
Reject	0/38 (0.0%)	0/10 (0.0%)		0/48 (0.0%)
Mean profit seller	72.16	62.40		70.13
Mean profit buyer	[24.68]	[37.60]		[27.38]
Mean total surplus	[96.84]	[100.00]		[97.50]
Low type:				
Buy A	2/34 (5.9%)**	2/14 (14.3%)**		4/48 (8.3%)**
Buy B	29/34 (85.3%)**			29/48 (60.4%)**
Reject	3/34 (8.8%)	12/14 (85.7%)**		15/48 (31.3%)**
Mean profit seller	26.32**	5.64**		20.29**
Mean profit buyer	[1.62**]	[0.07**]		[1.17**]
Mean total surplus	[27.94**]	[5.71**]		[21.46**]

Table 8. Private information computer treatment PIC-I. The stars indicate whether there are significant differences (* at the 5% level, ** at the 1% level) between the high and the low types.

	Menu	Only A		Total
# Obs. (Share)	33 (36.7%)	57 (63.3%)		90
Mean p_A	67.24	69.79		68.86
Mean p_B	32.79			32.79
High type:				
Buy A	13/14 (92.9%)	31/31 (100.0%)		44/45 (97.8%)
Buy B	1/14 (7.1%)			1/45 (2.2%)
Reject	0/14 (0.0%)	0/31 (0.0%)		0/45 (0.0%)
Mean profit seller	67.21	68.39		68.02
Mean profit buyer	[28.50]	[31.61]		[30.64]
Mean total surplus	[95.71]	[100.00]		[98.67]
Low type:				
Buy A	16/19 (84.2%)	23/26 (88.5%)		39/45 (86.7%)
Buy B	2/19 (10.5%)			2/45 (4.4%)
Reject	1/19 (5.3%)	3/26 (11.5%)		4/45 (8.9%)
Mean profit seller	52.58*	60.50		57.16*
Mean profit buyer	[9.53**]	[1.42**]		[4.84**]
Mean total surplus	[62.11**]	[61.92**]		[62.00**]

Table 9. Private information computer treatment PIC-II. The stars indicate whether there are significant differences (* at the 5% level, ** at the 1% level) between the high and the low types.

In PIC-I, the most frequently made offer was a menu with $p_A = 90$, $p_B = 30$, which is the theoretically optimal offer. This offer was made 14 times. There were 11 sellers who offered a menu with $p_A = 88$, $p_B = 29$, obviously because they were not absolutely sure that the computer would really act in their interest in case of indifference. Altogether, 33 sellers offered incentive-compatible menus with $p_A \in \{88, 89, 90\}$ and $p_B \in \{28, 29, 30\}$. Hence, 45.8% of all menus were roughly equal to the theoretically optimal solution. For comparison, in PI-I only one seller made such an offer.

In PIC-II, the most frequently observed decision of the sellers was to offer only good A at the price $p_A = 70$, which is the theoretically optimal offer. This offer was made 21 times. Altogether, 42 sellers offered only good A at a price $p_A \in \{68, 69, 70\}$. In addition, there were 13 sellers who offered equivalent menus. Thus, 61.1% of all sellers made an offer roughly equal to the theoretically optimal solution (whereas in PI-II, only 9.8% of the sellers made such offers).

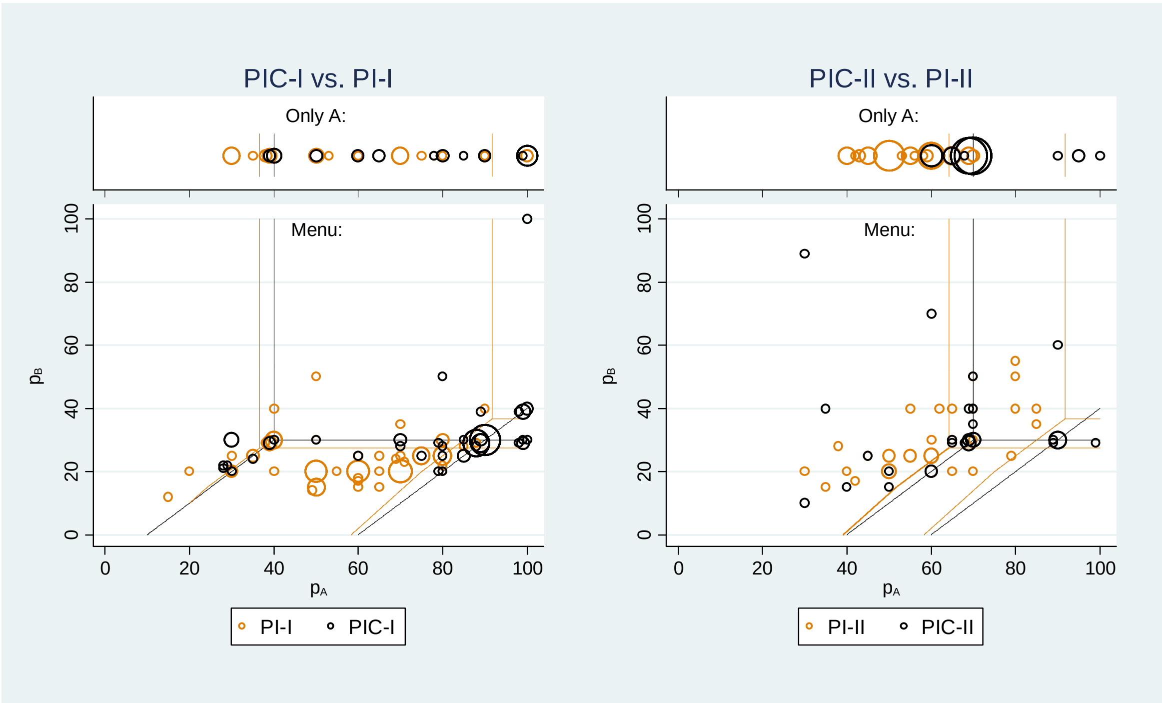

Figure 4. The distribution of the sellers' offers in the treatments PIC-I and PIC-II, compared to the offers in PI-I and PI-II. In each treatment, the size of the circles is proportional to the relative frequency with which the offers were made.

We can thus conclude that many sellers were able to find the optimal solution (doing so was clearly more difficult in PIC-I than in PIC-II), and inspection of Figure 4 shows that several other sellers were not too far away from the optimum.[22] Hence, a substantial amount of the deviations of the

[22]Not all of the noise that remains even in the computer treatments has to be attributed to decision errors. In particular, the sellers may be heterogenous with regard to their risk attitudes. Thus, at the end of each PIC-I session, sellers could choose one of three lotteries: (i) 100 ECU or 0 ECU with equal probability, (ii) 90 ECU or 30 ECU with equal probability, (iii) 40 ECU for sure. 80.2% of the sellers picked lottery (ii). Among the sellers whose contract offer roughly equaled the theoretical prediction, 93.9% chose lottery (ii), while 73.0% of the other sellers did so. At the end of PIC-II, sellers could choose lottery (i), (ii),

sellers' offers in PI-I and PI-II from the predictions based on standard theory cannot simply be attributed to decision errors. Instead, the sellers anticipate that the buyers will react as analyzed above and adjust their offers accordingly.

6 Conclusion

In a large-scale laboratory experiment, we have studied the simplest conceivable mechanism design problem that captures the main features of the basic adverse selection model, which plays a key role in the vast theoretical literature on contracts and incentives. A fundamental hypothesis of incentive theory is that a seller who faces a privately informed buyer may want to offer an incentive-compatible menu of contracts to separate the buyer types. In this case, the same total surplus level as under symmetric information will be attained in the good state of nature, while the presence of private information reduces the total surplus in the bad state of nature. Overall, our experimental results largely confirm the main predictions made by adverse selection theory.

However, in line with experimental studies of the ultimatum game, buyers tend to reject particularly inequitable offers, which may be explained by other-regarding preferences. While it turns out that the average strength of such preferences is relatively small, the fact that there is uncertainty about how a particular buyer of a given type will react to a specific offer is well anticipated by the sellers.

We also find that the presence of private information may increase the total surplus in the good state of nature, which might be surprising at first sight. Yet, the reason for this finding is that standard theory is too optimistic about the efficiency of contracting under symmetric information. Once we take into account that particularly inequitable offers are more likely to be rejected, the finding actually follows from adverse selection theory, which predicts that in the presence of private information sellers may set substantially smaller prices

or (iii') 70 ECU for sure. 80.0% of the sellers preferred lottery (iii'). 90.9% of the sellers whose contract offer roughly equaled the theoretical prediction chose lottery (iii'), while 62.9% of the other sellers did so. In any case, the vast majority of sellers preferred the lottery that is optimal under risk-neutrality. Thus, although risk preferences may explain some deviations from the predicted contract offers, they did not seem to play a major role.

than they would do in the good state of nature when there is symmetric information. Our experiment thus suggests that delineating the scope for efficiency *improvements* due to private information when buyers may be resentful could be an interesting topic for future theoretical research.[23]

Finally, in the contract-theoretic literature the basic adverse selection model has been extended in various directions. For instance, theorists have investigated the implications of verifiable ex post signals, type-dependent reservation utilities, and common values.[24] It could be a very promising avenue for future experimental research to explore some of these extensions in the laboratory.

[23] In the contract-theoretic literature, it is well-known that the presence of private information can be beneficial in the absence of full commitment power. For instance, in an incomplete contracting world, the hold-up problem may be ameliorated when there is private information (see e.g. Schmitz, 2006, and the literature discussed there). In contrast, our experimental findings show that private information may be beneficial even in a traditional mechanism design setting with full commitment.

[24] See Riordan and Sappington (1988) on how verifiable ex post signals can improve contracting. Lewis and Sappington (1989) study models with type-dependent reservation utilities, which may lead to countervailing incentives. Laffont and Martimort (2002) illustrate how common values may lead to non-responsiveness of the incentive scheme. See Nöldeke and Samuelson (2007) for a general adverse selection model that allows for common values.

Appendix

Proof of Proposition 2.

Let $x = (x^H, x^L)$ denote the trade decisions $x^\theta \in \{A, B, 0\}$ that the seller wants to implement for $\theta \in \{H, L\}$, where the decision 0 means no trade. Suppose first that the seller wants to implement the trade profile $x = (A, B)$. She maximizes her expected profit $\pi p_A + (1-\pi)p_B$ subject to the buyer's incentive compatibility constraints

$$v_A^H - p_A \geq v_B^H - p_B, \tag{IC$_H$}$$

$$v_B^L - p_B \geq v_A^L - p_A, \tag{IC$_L$}$$

and the buyer's participation constraints

$$v_A^H - p_A \geq 0, \tag{PC$_H$}$$

$$v_B^L - p_B \geq 0. \tag{PC$_L$}$$

Observe that (PC$_H$) is redundant, as it is implied by (IC$_H$) and (PC$_L$). Now ignore for a moment (IC$_L$), which will turn out to be satisfied by the solution. In the relaxed problem, (PC$_L$) must be binding, because otherwise the seller could increase her profit by increasing p_B without violating (IC$_H$). Hence, the seller sets $p_B = v_B^L$. Note that also (IC$_H$) must be binding, which implies $p_A = v_A^H - (v_B^H - v_B^L)$. It is straightforward to check that (IC$_L$) is satisfied given the Spence-Mirrlees condition $v_A^H - v_B^H \geq v_A^L - v_B^L$. The seller's expected profit if she implements $x = (A, B)$ is thus given by $\pi(v_A^H - (v_B^H - v_B^L)) + (1-\pi)v_B^L$.

Suppose now that the seller wants to implement the trade profile $x = (A, A)$. In this case, it is optimal for her to offer only good A at the price $p_A = v_A^L$. (Note that equivalently she could offer a menu with $p_A = v_A^L$ and $p_B \geq v_B^L$, since the Spence-Mirrlees condition is assumed to hold.) Thus, the seller's profit is v_A^L. Observe that the trade profile (A, B) yields a larger expected profit for the seller than (A, A) whenever $v_A^L < v_B^L + \pi(v_A^H - v_B^H)$.

Next, note that the trade profile (B, A) is not implementable due to the Spence-Mirrlees condition. It is also straightforward to show that the trade profiles $(0, A)$ and $(0, B)$ are not implementable. Moreover, it is easy to verify that under our assumptions the trade profiles (B, B), $(A, 0)$, $(B, 0)$, and $(0, 0)$ cannot be optimal for the seller. Proposition 2 follows immediately. □

References

ASPAROUHOVA, E. (2006), "Competition in Lending: Theory and Experiments", *Review of Finance*, **10**, 189–219.

BARON, D.P. and MYERSON, R.B. (1982), "Regulating a Monopolist with Unknown Costs", *Econometrica*, **50**, 911–930.

BLANCO, M., ENGELMANN, D. and NORMANN, H.T. (2011), "A Within-Subject Analysis of Other-Regarding Preferences", *Games and Economic Behavior*, **72**, 321–338.

BOLTON, P. and DEWATRIPONT, M. (2005), *Contract Theory* (Cambridge: MIT Press).

CABRALES, A., CHARNESS, G. and VILLEVAL, M.-C. (2011), "Hidden Information, Bargaining Power, and Efficiency: An Experiment", *Experimental Economics*, **14**, 133–159.

CAMERER, C.F. (2003), *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton: Princeton University Press).

CHARNESS, G. and DUFWENBERG, M. (2011), "Participation", *American Economic Review*, **101**, 1213–1239.

CHIAPPORI, P.A. and SALANIÉ, B. (2003), "Testing Contract Theory: A Survey of Some Recent Work", in Dewatripont, M., Hansen, L.P. and Turnovsky, S.J. (eds) *Advances in Economics and Econometrics,* Econometric Society Monographs, Eighth World Congress, Vol. 1 (Cambridge: Cambridge University Press) 115–149.

CRÉMER, J. and KHALIL, F. (1992), "Gathering Information before Signing a Contract", *American Economic Review*, **82**, 566–578.

CROSON, R.T.A. (1996), "Information in Ultimatum Games: An Experimental Study", *Journal of Economic Behavior and Organization*, **30**, 197–212.

DE BRUYN, A. and BOLTON, G.E. (2008), "Estimating the Influence of Fairness on Bargaining Behavior", *Management Science*, **54**, 1774–1791.

FEHR, E. and SCHMIDT, K.M. (1999), "A Theory of Fairness, Competition, and Cooperation", *Quarterly Journal of Economics*, **114**, 817–868.

FEHR, E. and SCHMIDT, K.M. (2006), "The Economics of Fairness, Reciprocity and Altruism - Experimental Evidence and New Theories", in Kolm, S.-C. and Ythier, J.M. (eds) *Handbook on the Economics of Giving, Reciprocity and Altruism,* Vol. 1 (Amsterdam: Elsevier) 615–691.

FISCHBACHER, U. (2007), "z-Tree: Zurich Toolbox for Ready-Made Economic Experiments", *Experimental Economics*, **10**, 171–178.

FORSYTHE, R., KENNAN, J. and SOPHER, B. (1991), "An Experimental Analysis of Strikes in Bargaining Games with One-Sided Private Information", *American Economic Review*, **81**, 253–278.

FUDENBERG, D. and TIROLE, J. (1991), *Game Theory*, Cambridge, MA: MIT Press.

GOLDMAN, M.B., LELAND, H.E., and SIBLEY, D.S. (1984), "Optimal Nonuniform Pricing", *Review of Economic Studies*, **51**, 305–320.

GREINER, B. (2004), "An Online Recruiting System for Economic Experiments", in Kremer, K. and Macho, V. (eds) *Forschung und wissenschaftliches Rechnen 2003. GWDG Bericht 63* (Göttingen: Ges. für Wiss. Datenverarbeitung) 79–93.

GUESNERIE, R. and LAFFONT, J.-J. (1984), "A Complete Solution to a Class of Principal-Agent Problems with an Application to the Control of a Self-Managed Firm", *Journal of Public Economics,* **25**, 329–369.

GÜTH, W. (1995), "On Ultimatum Bargaining Experiments – A Personal Review", *Journal of Economic Behavior and Organization,* **27**, 329–344.

GÜTH, W., HUCK, S. and OCKENFELS, P. (1996), "Two-Level Ultimatum Bargaining with Incomplete Information: An Experimental Study", *Economic Journal*, **106**, 593–604.

GÜTH, W., SCHMITTBERGER, R. and SCHWARZE, B. (1982), "An Experimental Analysis of Ultimatum Bargaining", *Journal of Economic Behavior and Organization*, **3**, 367–388.

GÜTH, W. and TIETZ, R. (1990), "Ultimatum Bargaining Behavior: A Survey and Comparison of Experimental Results", *Journal of Economic Psychology,* **11**, 417–449.

GÜTH, W. and VAN DAMME, E. (1998), "Information, Strategic Behavior, and Fairness in Ultimatum Bargaining: An Experimental Study", *Journal of Mathematical Psychology*, **42**, 227–247.

HARSTAD, R. and NAGEL, R. (2004), "Ultimatum Games with Incomplete Information on the Side of the Proposer: An Experimental Study", *Cuadernos de Economía*, **27**, 37–74.

HART, O. and HOLMSTRÖM, B. (1987), "The Theory of Contracts", in Bewley, T. (eds) *Advances in Economics and Econometrics,* Econometric Society Monographs, Fifth World Congress (Cambridge: Cambridge University Press) 71–155.

HOFFMAN, E., MCCABE, K., SHACHAT, K. and SMITH, V. (1994), "Preferences, Property Rights, and Anonymity in Bargaining Games", *Games and Economic Behavior*, **7**, 346–380.

HOPPE, E.I. and SCHMITZ, P.W. (2013), "Contracting under Incomplete Information and Social Preferences: An Experimental Study", *Review of Economic Studies*, forthcoming.

HOUSER, D. and KURZBAN, R. (2002), "Revisiting Kindness and Confusion in Public Goods Experiments", *American Economic Review*, **92**, 1062–1069.

HUCK, S. (1999), "Responder Behavior in Ultimatum Offer Games with Incomplete Information", *Journal of Economic Psychology*, **20**, 183–206.

HUCK, S., SELTZER, A.J. and WALLACE, B. (2011), "Deferred Compensation in Multiperiod Labor Contracts: An Experimental Test of Lazear's Model", *American Economic Review*, **101**, 819–843.

KAGEL, J.H., KIM, C. and MOSER, D. (1996), "Fairness in Ultimatum Games with Asymmetric Information and Asymmetric Payoffs", *Games and Economic Behavior*, **13**, 100–110.

KESSLER, A.S. (1998), "The Value of Ignorance", *Rand Journal of Economics*, **29**, 339–354.

LAFFONT, J.-J. and MARTIMORT, D. (2002), *The Theory of Incentives: The Principal-Agent Model* (Princeton: Princeton University Press).

LAFFONT, J.-J. and TIROLE, J. (1993), *A Theory of Incentives in Procurement and Regulation* (Cambridge: MIT Press).

LEWIS, T.R. and SAPPINGTON, D.E.M. (1989), "Countervailing Incentives in Agency Problems", *Journal of Economic Theory*, **49**, 294–313.

LEWIS, T.R. and SAPPINGTON, D.E.M. (1997), "Information Management in Incentive Problems", *Journal of Political Economy*, **105**, 796–821.

MASKIN, E. and RILEY, J. (1984), "Monopoly with Incomplete Information", *Rand Journal of Economics*, **15**, 171–196.

MCKELVEY, R.D. and PALFREY, T.R. (1995), "Quantal Response Equilibria in Normal Form Games", *Games and Economic Behavior*, **10**, 6–38.

MCKELVEY, R.D. and PALFREY, T.R. (1998), "Quantal Response Equilibria in Extensive Form Games", *Experimental Economics*, **1**, 9–41.

MIRRLEES, J.A. (1971), "An Exploration in the Theory of Optimum Income Taxation", *Review of Economic Studies*, **38**, 175–208.

MITZKEWITZ, M. and NAGEL, R. (1993), "Experimental Results on Ultimatum Games with Incomplete Information", *International Journal of Game Theory*, **22**, 171–198.

MUSSA, M. and ROSEN, S. (1978), "Monopoly and Product Quality", *Journal of Economic Theory*, **18**, 301–317.

NOBEL PRIZE COMMITTEE (2007), "Mechanism Design Theory", The Royal Swedish Academy of Sciences.

NÖLDEKE, G. and SAMUELSON, L. (2007), "Optimal Bunching Without Optimal Control", *Journal of Economic Theory*, **134**, 405–420.

PRENDERGAST, C. (1999), "The Provision of Incentives in Firms", *Journal of Economic Literature*, **37**, 7–63.

RAPOPORT, A. and SUNDALI, J.A. (1996), "Ultimatums in Two-Person Bargaining with One-Sided Uncertainty: Offer Games", *International Journal of Game Theory*, **25**, 475–494.

RIORDAN, M.H. and SAPPINGTON, D.E. (1988), "Optimal Contracts with Public Ex Post Information", *Journal of Economic Theory*, **45**, 189–199.

ROGERS, B.W., PALFREY, T.R. and CAMERER, C.F. (2009), "Heterogeneous Quantal Response Equilibrium and Cognitive Hierarchies", *Journal of Economic Theory*, **144**, 1440–1467.

SALANIÉ, B. (2005), *The Economics of Contracts* (Cambridge: MIT Press).

SCHMITZ, P.W. (2006), "Information Gathering, Transaction Costs, and the Property Rights Approach", *American Economic Review*, **96**, 422–434.

STRAUB, P.G. and MURNIGHAN, J.K. (1995), "An Experimental Investigation of Ultimatum Games: Information, Fairness, Expectations, and Lowest Acceptable Offers", *Journal of Economic Behavior and Organization*, **27**, 345–364.

Supplementary Material

for

**"Do Sellers Offer Menus of Contracts to Separate Buyer Types?
An Experimental Test of Adverse Selection Theory"**

(Eva I. Hoppe and Patrick W. Schmitz)

The *Supplementary Material* contains the instructions that were handed out to the participants in the six treatments of our experiment.

Instructions for the PI-I treatment:

Experimental Instructions

In this experiment there is always one seller who interacts with one buyer. You will be randomly assigned either to the role of the seller or to the role of the buyer.

The currency in the experiment is called ECU (Experimental Currency Unit).

The seller can sell **either** good A **or** good B to the buyer.

There are two types of buyers, which have the same probability:

- A buyer of type 1 has a valuation of 100 ECU for good A and a valuation of 40 ECU for good B.
- A buyer of type 2 has a valuation of 40 ECU for good A and a valuation of 30 ECU for good B.

The buyer knows his type. The seller only knows that the buyer with whom he is matched is either of type 1 or of type 2, with 50% probability each.

The seller has no costs in this experiment. This means that if he sells a good, then the obtained price is his profit.

In detail, the experiment proceeds as follows:

The experiment consists of only one single period.
The period consists of two stages:

Stage 1: Seller makes offer

On the screen you can see whether you have been assigned to the role of the seller or to the role of the buyer. **Only the buyer** learns whether he is of type 1 or of type 2.

The seller decides first how he wants to design his offer for sale. He can choose one of three alternatives:
1) He offers only good A to the buyer.
2) He offers only good B to the buyer.
3) He offers both goods; the buyer can buy only one of them.

After that, the seller sets the price(s):

If the seller has chosen alternative 1, he sets a price p_A for good A.
If the seller has chosen alternative 2, he sets a price p_B for good B.
If the seller has chosen alternative 3, he sets a price p_A for good A and a price p_B for good B.

(A price has to be an integer between 0 ECU and 100 ECU.)

Stage 2: Buyer makes buying decision

The buyer learns the seller's offer.

If the seller has offered only one good for sale, then the buyer decides whether or not he buys this good.

If the seller has offered good A and good B, then the buyer decides whether he buys good A, good B, or no good.

The profits are as follows.

Seller's profit (in ECU):

If the buyer has bought good A:	p_A
If the buyer has bought good B:	p_B
If the buyer has bought no good:	0

Buyer's profit (in ECU):

	If the buyer is of type 1:	If the buyer is of type 2:
If the buyer has bought good A:	$100 - p_A$	$40 - p_A$
If the buyer has bought good B:	$40 - p_B$	$30 - p_B$
If the buyer has bought no good:	0	0

Your payoff:
In addition to the (possibly negative) profit realized in the experiment you get 70 ECU and the resulting amount will be paid out to you in cash at an exchange rate of 0.12 Euro per ECU.

Please note:
During the whole experiment communication is not allowed. If you have a question, please raise your hand out of the cabin. All decisions are anonymous; i.e., no participant ever learns the identity of a person who has made a particular decision. The payment is conducted anonymously, too; i.e., no participant learns what the payoff of another participant is.

Instructions for the PI-II treatment:

Experimental Instructions

In this experiment there is always one seller who interacts with one buyer. You will be randomly assigned either to the role of the seller or to the role of the buyer.

The currency in the experiment is called ECU (Experimental Currency Unit).

The seller can sell **either** good A **or** good B to the buyer.

There are two types of buyers, which have the same probability:

- A buyer of type 1 has a valuation of 100 ECU for good A and a valuation of 40 ECU for good B.
- A buyer of type 2 has a valuation of 70 ECU for good A and a valuation of 30 ECU for good B.

The buyer knows his type. The seller only knows that the buyer with whom he is matched is either of type 1 or of type 2, with 50% probability each.

The seller has no costs in this experiment. This means that if he sells a good, then the obtained price is his profit.

In detail, the experiment proceeds as follows:

The experiment consists of only one single period.
The period consists of two stages:

Stage 1: Seller makes offer

On the screen you can see whether you have been assigned to the role of the seller or to the role of the buyer. **Only the buyer** learns whether he is of type 1 or of type 2.

The seller decides first how he wants to design his offer for sale. He can choose one of three alternatives:
1) He offers only good A to the buyer.
2) He offers only good B to the buyer.
3) He offers both goods; the buyer can buy only one of them.

After that, the seller sets the price(s):

If the seller has chosen alternative 1, he sets a price p_A for good A.
If the seller has chosen alternative 2, he sets a price p_B for good B.
If the seller has chosen alternative 3, he sets a price p_A for good A and a price p_B for good B.

(A price has to be an integer between 0 ECU and 100 ECU.)

Stage 2: Buyer makes buying decision

The buyer learns the seller's offer.

If the seller has offered only one good for sale, then the buyer decides whether or not he buys this good.

If the seller has offered good A and good B, then the buyer decides whether he buys good A, good B, or no good.

The profits are as follows.

Seller's profit (in ECU):

If the buyer has bought good A:	p_A
If the buyer has bought good B:	p_B
If the buyer has bought no good:	0

Buyer's profit (in ECU):

	If the buyer is of type 1:	If the buyer is of type 2:
If the buyer has bought good A:	$100 - p_A$	$70 - p_A$
If the buyer has bought good B:	$40 - p_B$	$30 - p_B$
If the buyer has bought no good:	0	0

Your payoff:
In addition to the (possibly negative) profit realized in the experiment you get 70 ECU and the resulting amount will be paid out to you in cash at an exchange rate of 0.12 Euro per ECU.

Please note:
During the whole experiment communication is not allowed. If you have a question, please raise your hand out of the cabin. All decisions are anonymous; i.e., no participant ever learns the identity of a person who has made a particular decision. The payment is conducted anonymously, too; i.e., no participant learns what the payoff of another participant is.

Instructions for the SI-I treatment:

Experimental Instructions

In this experiment there is always one seller who interacts with one buyer. You will be randomly assigned either to the role of the seller or to the role of the buyer.

The currency in the experiment is called ECU (Experimental Currency Unit).

The seller can sell **either** good A **or** good B to the buyer.

There are two types of buyers, which have the same probability:

- A buyer of type 1 has a valuation of 100 ECU for good A and a valuation of 40 ECU for good B.
- A buyer of type 2 has a valuation of 40 ECU for good A and a valuation of 30 ECU for good B.

Both the seller and the buyer know whether the buyer is of type 1 or of type 2.

The seller has no costs in this experiment. This means that if he sells a good, then the obtained price is his profit.

In detail, the experiment proceeds as follows:

The experiment consists of only one single period.
The period consists of two stages:

Stage 1: Seller makes offer

On the screen you can see whether you have been assigned to the role of the seller or to the role of the buyer. **Both the seller and the buyer** learn whether the buyer is of type 1 or of type 2.

The seller decides first how he wants to design his offer for sale. He can choose one of three alternatives:
1) He offers only good A to the buyer.
2) He offers only good B to the buyer.
3) He offers both goods; the buyer can buy only one of them.

After that, the seller sets the price(s):

If the seller has chosen alternative 1, he sets a price p_A for good A.
If the seller has chosen alternative 2, he sets a price p_B for good B.
If the seller has chosen alternative 3, he sets a price p_A for good A and a price p_B for good B.

(A price has to be an integer between 0 ECU and 100 ECU.)

Stage 2: Buyer makes buying decision

The buyer learns the seller's offer.

If the seller has offered only one good for sale, then the buyer decides whether or not he buys this good.

If the seller has offered good A and good B, then the buyer decides whether he buys good A, good B, or no good.

The profits are as follows.

Seller's profit (in ECU):

If the buyer has bought good A:	p_A
If the buyer has bought good B:	p_B
If the buyer has bought no good:	0

Buyer's profit (in ECU):

	If the buyer is of type 1:	If the buyer is of type 2:
If the buyer has bought good A:	$100 - p_A$	$40 - p_A$
If the buyer has bought good B:	$40 - p_B$	$30 - p_B$
If the buyer has bought no good:	0	0

Your payoff:
In addition to the (possibly negative) profit realized in the experiment you get 70 ECU and the resulting amount will be paid out to you in cash at an exchange rate of 0.12 Euro per ECU.

Please note:
During the whole experiment communication is not allowed. If you have a question, please raise your hand out of the cabin. All decisions are anonymous; i.e., no participant ever learns the identity of a person who has made a particular decision. The payment is conducted anonymously, too; i.e., no participant learns what the payoff of another participant is.

Instructions for the SI-II treatment:

Experimental Instructions

In this experiment there is always one seller who interacts with one buyer. You will be randomly assigned either to the role of the seller or to the role of the buyer.

The currency in the experiment is called ECU (Experimental Currency Unit).

The seller can sell **either** good A **or** good B to the buyer.

There are two types of buyers, which have the same probability:

- A buyer of type 1 has a valuation of 100 ECU for good A and a valuation of 40 ECU for good B.
- A buyer of type 2 has a valuation of 70 ECU for good A and a valuation of 30 ECU for good B.

Both the seller and the buyer know whether the buyer is of type 1 or of type 2.

The seller has no costs in this experiment. This means that if he sells a good, then the obtained price is his profit.

In detail, the experiment proceeds as follows:

The experiment consists of only one single period.
The period consists of two stages:

Stage 1: Seller makes offer

On the screen you can see whether you have been assigned to the role of the seller or to the role of the buyer. **Both the seller and the buyer** learn whether the buyer is of type 1 or of type 2.

The seller decides first how he wants to design his offer for sale. He can choose one of three alternatives:
1) He offers only good A to the buyer.
2) He offers only good B to the buyer.
3) He offers both goods; the buyer can buy only one of them.

After that, the seller sets the price(s):

If the seller has chosen alternative 1, he sets a price p_A for good A.
If the seller has chosen alternative 2, he sets a price p_B for good B.
If the seller has chosen alternative 3, he sets a price p_A for good A and a price p_B for good B.

(A price has to be an integer between 0 ECU and 100 ECU.)

Stage 2: Buyer makes buying decision

The buyer learns the seller's offer.

If the seller has offered only one good for sale, then the buyer decides whether or not he buys this good.

If the seller has offered good A and good B, then the buyer decides whether he buys good A, good B, or no good.

The profits are as follows.

Seller's profit (in ECU):

If the buyer has bought good A:	p_A
If the buyer has bought good B:	p_B
If the buyer has bought no good:	0

Buyer's profit (in ECU):

	If the buyer is of type 1:	If the buyer is of type 2:
If the buyer has bought good A:	$100 - p_A$	$70 - p_A$
If the buyer has bought good B:	$40 - p_B$	$30 - p_B$
If the buyer has bought no good:	0	0

Your payoff:
In addition to the (possibly negative) profit realized in the experiment you get 70 ECU and the resulting amount will be paid out to you in cash at an exchange rate of 0.12 Euro per ECU.

Please note:
During the whole experiment communication is not allowed. If you have a question, please raise your hand out of the cabin. All decisions are anonymous; i.e., no participant ever learns the identity of a person who has made a particular decision. The payment is conducted anonymously, too; i.e., no participant learns what the payoff of another participant is.

Instructions for the PIC-I treatment:

Experimental Instructions

In this experiment every participant is assigned to the role of a seller. You do *not* interact with any other participant of the experiment. You interact only with the computer. The computer takes on the role of the buyer.

The currency in the experiment is called ECU (Experimental Currency Unit).

You can sell **either** good A **or** good B to the computer.

The computer takes on either the role of a buyer of type 1 or the role of a buyer of type 2. Both types have the same probability:

- A buyer of type 1 has a valuation of 100 ECU for good A and a valuation of 40 ECU for good B.
- A buyer of type 2 has a valuation of 40 ECU for good A and a valuation of 30 ECU for good B.

The computer knows whether he is a buyer of type 1 or a buyer of type 2.
You only know that the computer takes on the role of a buyer of type 1 or the role of a buyer of type 2 with 50% probability each.

The seller has no costs in this experiment. This means that if you sell a good, then the obtained price is your profit.

In detail, the experiment proceeds as follows:

The experiment consists of only one single period.
The period consists of two stages:

Stage 1: Seller makes offer

You decide first how you want to design you offer for sale. You can choose one of three alternatives:
1) You offer only good A to the computer.
2) You offer only good B to the computer.
3) You offer both goods; the computer can buy only one of them.

After that, you set the price(s):

If you have chosen alternative 1, you set a price p_A for good A.
If you have chosen alternative 2, you set a price p_B for good B.
If you have chosen alternative 3, you set a price p_A for good A and a price p_B for good B.

(A price has to be an integer between 0 ECU and 100 ECU.)

Stage 2: Computer makes buying decision

If you have offered only one good for sale, then the computer decides whether or not he buys this good.

If you have offered good A and good B, then the computer decides whether he buys good A, good B, or no good.

The computer makes his buying decision such that his profit becomes as large as possible.
If in doing so the computer can make the same profit by different decisions, he chooses the decision that is best for you.

The profits are as follows.

Your profit (in ECU):

If the computer has bought good A:	p_A
If the computer has bought good B:	p_B
If the computer has bought no good:	0

Profit of the computer (in ECU):

	If the computer is a buyer of type 1:	If the computer is a buyer of type 2:
If the computer has bought good A:	$100 - p_A$	$40 - p_A$
If the computer has bought good B:	$40 - p_B$	$30 - p_B$
If the computer has bought no good:	0	0

Your payoff:
The profit that you have realized in the experiment will be paid out to you in cash at an exchange rate of 0.12 Euro per ECU.

Please note:
During the whole experiment communication is not allowed. If you have a question, please raise your hand out of the cabin. All decisions are anonymous; i.e., no participant ever learns the identity of a person who has made a particular decision. The payment is conducted anonymously, too; i.e., no participant learns what the payoff of another participant is.

<u>**Instructions for the PIC-II treatment:**</u>

Experimental Instructions

In this experiment every participant is assigned to the role of a seller. You do *not* interact with any other participant of the experiment. You interact only with the computer. The computer takes on the role of the buyer.

The currency in the experiment is called ECU (Experimental Currency Unit).

You can sell **either** good A **or** good B to the computer.

The computer takes on either the role of a buyer of type 1 or the role of a buyer of type 2. Both types have the same probability:

- A buyer of type 1 has a valuation of 100 ECU for good A and a valuation of 40 ECU for good B.
- A buyer of type 2 has a valuation of 70 ECU for good A and a valuation of 30 ECU for good B.

The computer knows whether he is a buyer of type 1 or a buyer of type 2.
You only know that the computer takes on the role of a buyer of type 1 or the role of a buyer of type 2 with 50% probability each.

The seller has no costs in this experiment. This means that if you sell a good, then the obtained price is your profit.

In detail, the experiment proceeds as follows:

The experiment consists of only one single period.
The period consists of two stages:

Stage 1: Seller makes offer

You decide first how you want to design you offer for sale. You can choose one of three alternatives:
1) You offer only good A to the computer.
2) You offer only good B to the computer.
3) You offer both goods; the computer can buy only one of them.

After that, you set the price(s):

If you have chosen alternative 1, you set a price p_A for good A.
If you have chosen alternative 2, you set a price p_B for good B.
If you have chosen alternative 3, you set a price p_A for good A and a price p_B for good B.

(A price has to be an integer between 0 ECU and 100 ECU.)

Stage 2: Computer makes buying decision

If you have offered only one good for sale, then the computer decides whether or not he buys this good.

If you have offered good A and good B, then the computer decides whether he buys good A, good B, or no good.

The computer makes his buying decision such that his profit becomes as large as possible.
If in doing so the computer can make the same profit by different decisions, he chooses the decision that is best for you.

The profits are as follows.

Your profit (in ECU):

If the computer has bought good A:	p_A
If the computer has bought good B:	p_B
If the computer has bought no good:	0

Profit of the computer (in ECU):

	If the computer is a buyer of type 1:	If the computer is a buyer of type 2:
If the computer has bought good A:	$100 - p_A$	$70 - p_A$
If the computer has bought good B:	$40 - p_B$	$30 - p_B$
If the computer has bought no good:	0	0

Your payoff:
The profit that you have realized in the experiment will be paid out to you in cash at an exchange rate of 0.12 Euro per ECU.

Please note:
During the whole experiment communication is not allowed. If you have a question, please raise your hand out of the cabin. All decisions are anonymous; i.e., no participant ever learns the identity of a person who has made a particular decision. The payment is conducted anonymously, too; i.e., no participant learns what the payoff of another participant is.