CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The University of Nottingham

UNITED KINGDOM · CHINA · MALAYSIA

Discussion Paper No. 2018-03

Alex Possajennikov and Rene Saran

May 2018

Inefficiency in Private Value Bargaining with Naive Players: An Experimental Study

CeDEx Discussion Paper Series
ISSN 1749 - 3293

CENTRE FOR DECISION RESEARCH & EXPERIMENTAL ECONOMICS

The Centre for Decision Research and Experimental Economics was founded in 2000, and is based in the School of Economics at the University of Nottingham.

The focus for the Centre is research into individual and strategic decision-making using a combination of theoretical and experimental methods. On the theory side, members of the Centre investigate individual choice under uncertainty, cooperative and non-cooperative game theory, as well as theories of psychology, bounded rationality and evolutionary game theory. Members of the Centre have applied experimental methods in the fields of public economics, individual choice under risk and uncertainty, strategic interaction, and the performance of auctions, markets and other economic institutions. Much of the Centre's research involves collaborative projects with researchers from other departments in the UK and overseas.

Please visit http://www.nottingham.ac.uk/cedex for more information about the Centre or contact

Suzanne Robey
Centre for Decision Research and Experimental Economics
School of Economics
University of Nottingham
University Park
Nottingham
NG7 2RD
Tel: +44 (0)115 95 14763
suzanne.robey@nottingham.ac.uk

The full list of CeDEx Discussion Papers is available at

http://www.nottingham.ac.uk/cedex/publications/discussion-papers/index.aspx

Inefficiency in Private Value Bargaining with Naive Players: An Experimental Study*

Alex Possajennikov[†]
University of Nottingham

Rene Saran
University of Cincinnati

May 2018

Abstract

The paper reports on an experiment on two-player double-auction bargaining with private values. We consider a setting with discrete two-point overlapping distributions of traders' valuations, in which there exists a fully efficient equilibrium. We show that if there are traders that behave naively, i.e., set bid or ask equal to their valuation, then there is no equilibrium achieving full efficiency. In the experiment, we vary the proportion of naive traders by introducing computerized players. We find that full efficiency is not achieved in the experiment with or without naive traders, and efficiency is not lower in the presence of naive traders. Subjects mostly set bid/ask prices strategically but the extent of strategic behavior is not larger in the presence of naive players. We can explain these results by a learning model of noisy strategy adjustment. We also find that framing the double auction as a direct mechanism leads to more naive behavior by experiment participants, and that allowing face-to-face pre-play communication increases efficiency although still not to the full level.

Keywords: bargaining with private values, double auction, efficiency, honesty

JEL Codes: C72, C78, C91, D82

*We thank the British Academy for financial support, Centre for Decision Research and Experimental Economics (CeDEx) at the University of Nottingham for access to laboratory facilities, and Lu Dong, Francesco Fallucchi, Lingbo Huang, Jörg Weber and Monika Pompeo for research assistance. The paper benefited from audience comments at seminar presentations at the University of the Basque Country (Bilbao), University of Mannheim and Tilburg University, as well as conference presentations at the International Meeting on Exprerimental and Behavioral Social Sciences (IMEBESS 2014), 2015 Lancaster Game Theory conference, 5th World Meeting of the Game Theory Society (GAMES 2016), and 2017 Royal Economic Society (RES) conference.

[†]Corresponding author: School of Economics, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom. Tel: +44 115 9515461, email: alex.possajennikov@nottingham.ac.uk, http://www.nottingham.ac.uk/~lezap.

1 Introduction

Many goods, e.g., cars, homes, and novelty items in a bazaar are traded in bargaining encounters involving one buyer and one seller. The buyer and seller are often asymmetrically informed: each trader knows his/her own reservation price (value or cost) but not that of the other trader. In such private value bargaining negotiations, a trader faces the basic tradeoff between the probability and terms of trade: misrepresenting one's reservation price can improve the terms of trade but can reduce the probability of trade. Strategic traders optimally respond to this tradeoff, which can result in an inefficient outcome, i.e., the good is not traded even though the buyer's value exceeds the seller's cost (Myerson and Satterthwaite, 1983).

Experimental evidence, however, suggests that some people behave naively, i.e., they do not necessarily misrepresent their private information even though they might have monetary incentives to do so.[1] Saran (2011) incorporates the possibility of naive traders in private value bargaining to show that a mechanism designer who tailors the bargaining mechanism to the proportion of naive traders can improve efficiency. However, if the mechanism is fixed, then the introduction of naive traders can increase or decrease efficiency because, depending on the mechanism, strategic traders could react less or more strategically to the presence of naive traders.[2]

In this paper, we experimentally investigate the effect of introducing naive traders on efficiency and individual behavior in private values bargaining. We are specially interested in testing the theoretical possibility that the presence of naive traders can reduce efficiency in this environment. To that end, we consider the double-auction mechanism and a simple setting with discrete (two points) overlapping but non-identical distributions of values and costs. In the double auction, the buyer submits a bid while the seller submits an ask and trade happens if and only if the bid weakly exceeds the ask, at the price midpoint between the bid and the ask. Without naive traders, there exists a fully efficient equilibrium in our

[1] See, for example, Fischbacher and Föllmi-Heusi (2013), Abeler et al. (2016) and Gneezy et al. (2018) for evidence of this in simple one-person reporting decisions, and Gneezy (2005), Lundquist et al. (2009) and Serra-Garcia et al. (2011) for interactive situations. These studies suggest aversion to lying (and aversion to being seen lying) as an explanation of naive behavior. Bounded rationality, i.e., inability to understand the implications of revealing one's private information could also explain naive behavior in some circumstances. For example, although people say that they are concerned about privacy, they willingly reveal personal information on the internet (e.g., Spiekermann et al., 2001).

[2] For example, Saran (2012) shows that the presence of naive traders in a double auction with pre-play communication can improve efficiency since the strategic traders will act less strategically in the pre-play communication stage lest they lose the chance to trade with the naive traders in the double-auction stage at a favorable price. Without pre-play communication, Saran (2011) shows that the presence of naive traders in a double auction can reduce efficiency when the intervals of values and costs overlap at only one point.

setting. However, as the proportion of naive traders increases, strategic traders have an incentive to increase their strategic behavior and the efficient equilibria disappear.[3] Thus, if the traders coordinate on the most efficient equilibrium, then the introduction of naive traders can reduce efficiency in this setting in theory.

Our confidence in this theoretical possibility is buttressed by previous experimental studies on the double auction which show that the outcomes are consistent with the most efficient equilibrium. For the case of identical continuous uniform distributions of values and costs, the linear equilibrium of the double auction suggested by Chatterjee and Samuelson (1983) is the most efficient (even though not all trades with positive surplus are achieved in it).[4] Although there are many other equilibria of the double auction (Leininger et al., 1989; Satterthwaite and Williams, 1989), experimental data typically conform to the theoretical predictions of the linear equilibrium (Radner and Schotter, 1989; Valley et al., 2002; McGinn et al., 2003; Ellingsen et al., 2009).

We run two main types of sessions in our experiment. In one set of sessions, each subject always plays against another subject in the experiment, with one subject assigned the role of the buyer while the other the role of the seller. In another set of sessions, with some probability, a subject, instead of being matched with another subject, is matched with a computer that will set ask/bid equal to cost/value. Subjects know this possibility but at the time of making their decisions, they do not know whether they will be matched with a computer or with another subject. In this way, the chance of playing a naive opponent increases with the introduction of the artificial naive subjects, and we can analyze the effect of this on the behavior of human participants.

We also compare the performance of the double-auction mechanism along another dimension, namely whether framing the mechanism as a bid-ask setting, or as a direct mechanism with the same allocation rule makes a difference.[5] Theoretically this does not affect the set of equilibria, but the direct mechanism may be more conducive to naive play if subjects are averse to lying since it is clear what telling the truth means there.

The experiment results are mixed, with no clear effect of the change in the proportion of naive players on efficiency, or on other variables such as the proportion of naive plays by human participants or the extent of setting ask/bid different from cost/value. The subjects

[3] There exist other (including mixed-strategy) equilibria in our setting with or without naive traders but none of them can achieve full efficiency.

[4] Myerson and Satterthwaite (1983) show that the linear equilibrium is in fact constrained efficient for the case of identical continuous uniform distributions, i.e., there is no other bargaining mechanism which would lead in equilibrium to a higher ex-ante expected gains from trade.

[5] In the direct mechanism, players make direct statements about their values/costs. Trade happens if the buyer's statement exceeds the seller's, at the price equal to the midpoint between the stated value and cost.

do not coordinate on the fully efficient equilibrium in the setting without artificial naive players. The increase in the proportion of naive players does not decrease efficiency in the experiment (efficiency in fact increases for some settings but this is solely due to the positive effect of artificial naive traders rather than due to a change in strategic behavior by human subjects). The direct mechanism frame has an effect to make some traders behave naively more often, but this effect is not sufficient to change efficiency.

We also look at the patterns of subjects' individual decisions and try to identify what determines whether the play is naive and the extent of strategic shading or exaggeration. We find that subjects do respond to incentives to some extent: the incidence of naive play is lower if the payoff obtained from lying is higher, and the bids and asks adjust towards the best response to the observed history. Most subjects thus do behave strategically, but the extent of strategic behavior does not change significantly across different settings to influence the aggregate measures of efficiency.

That the observed history significantly affects individual behavior suggests that, instead of equilibrium, a learning model might be more appropriate to model behavior. We specifically consider the experienced-weighted adjustment (EWA) learning model of noisy best response (Camerer and Ho, 1999). The model is initialized with behavior based on low levels of strategic sophistication in the level-k model (e.g., Stahl and Wilson, 1994)[6], which turns out to resemble the initial distribution of choices.[7] For a reasonable set of parameters, simulations of the model reproduce most dynamic features of the experiment results, also finding that efficiency and the extent of strategic behavior do not change much with a change in the proportion of naive players.

In addition to the main set of experimental sessions on the double auction without communication, we also run sessions that include a face-to-face pre-play communication stage but without artificial naive players. In line with previous experiments (Valley et al., 2002; McGinn et al., 2003), we find that efficiency is higher with pre-play communication than without it mainly due to two "dyadic" strategies: (a) mutual revelation of value and cost and (b) coordination on a single price. However, in contrast to the conclusion of these previous experiments that pre-play communication can generate efficiency gains that breach the theoretical maximum, we do not find subjects attaining full efficiency even with pre-play communication although it is theoretically feasible.[8] The reason for this result in our

[6]More precisely, the initial attractions of the EWA model are either set equal to the expected payoff against the uniform distribution of strategies of the other player (level-1) or to the expected payoffs against noisy best response to such a uniform distribution (level-2).

[7]This adds to the growing evidence that the level-k model is a better predictor of initial responses in games than equilibrium models (see Crawford et al., 2013 for a survey). Crawford (2016) and Kneeland (2017) apply the level-k theory to private value bargaining from the point of view of mechanism design.

[8]The previous experiments used the setting with identical continuous uniform distributions of values and

setting is that there is a non-negligible fraction of traders who misrepresent their cost or value, sometimes leading to an impasse in trading even though gains from trade are positive.

To sum up, although the efficient outcome could be achieved in equilibrium in our setting, such an outcome is not obtained in the experiment, even with pre-play communication. Failure to attain full efficiency in our experiments cautions against optimistic predictions based on the analysis of equilibrium efficiency properties. Although the previous experiments on the double auction without communication (Radner and Schotter, 1989; Valley et al., 2002; McGinn et al., 2003; Ellingsen et al., 2009) found that players' behavior appears to be consistent with the (constrained efficient) linear equilibrium, our findings suggest that play is better described by low levels of strategic sophistication. These two seemingly contradictory observations can be reconciled by noting that the linear equilibrium of the double auction in the setting with identical uniform distributions of values is also consistent with level-k play.[9] Thus, non-equilibrium models like level-k and learning dynamics may be more useful predictors of play in the double auction.

2 Double Auction Bargaining with Discrete Private Values

2.1 Model

There are two risk-neutral traders, a single seller (s) of an indivisible object facing a single buyer (b). The seller's cost c of producing the object can be either low $\underline{c}$ of high $\bar{c}$ with $q_s = \frac{1}{2}$ being the probability of $\underline{c}$. Similarly, the buyer's value v for the object can be either low $\underline{v}$ or high $\bar{v}$ with $q_b = \frac{1}{2}$ being the probability of $\bar{v}$. The seller privately knows her cost while the buyer privately knows her value. We assume $\underline{c} < \underline{v} < \bar{c} < \bar{v}$.[10]

The traders use the double auction as the trading mechanism. In this mechanism, the two traders simultaneously submit a bid (buyer) and an ask (seller). If the buyer's bid z_b weakly exceeds the seller's ask z_s, then the object is traded at price $p = \frac{1}{2}(z_b + z_s)$. Then the buyer's payoff is $v - p$, while the seller's payoff is $p - c$. On the other hand, if $z_b < z_s$, then there is no trade and each trader obtains a zero payoff.

Each trader i can have one of two dispositions t_i: strategic (str) or naive (n). In the double auction, the naive buyer sets bid equal to her value while the naive seller sets ask

costs. They found that subjects were able to capture gains from trade that were greater than those in the (constrained efficient) linear equilibrium.

[9] For instance, the linear equilibrium strategy arises in expectation if there are 25% of (non-noisy) level-1 players and 75% of (non-noisy) level-2 players. See Crawford (2016) for the details of level-1 and level-2 strategies in this setting.

[10] Chatterjee and Samuelson (1987) consider such a setting for a model of alternating offer bargaining; Feri and Gantner (2011) run an experiment based on their model.

equal to her cost. The strategic trader can set any bid/ask in the double auction. The traders do not know each other's disposition. The probability that a trader is naive is $\varepsilon \in [0, 1)$, which is independent of the trader's value/cost.

2.2 Equilibrium and Efficiency

Let $Z_s^{t_s} : \{\underline{c}, \bar{c}\} \to \mathbb{R}$ and $Z_b^{t_b} : \{\underline{v}, \bar{v}\} \to \mathbb{R}$ denote the pure strategies of, respectively, the seller and buyer in the double auction. For the naive traders, by assumption $Z_s^n(c) = c$ and $Z_b^n(v) = v$. Thus we only need to analyze the strategies of strategic traders Z_c^{str} and Z_b^{str}.

Let $q^{t_s,t_b}(\cdot,\cdot) : \{\underline{c}, \bar{c}\} \times \{\underline{v}, \bar{v}\} \to [0, 1]$ denote the probability of trade between the traders depending on their dispositions and values. Given the assumption that $\underline{c} < \underline{v} < \bar{c} < \bar{v}$, the efficient trade occurs if the seller and the buyer trade in all cases except when cost is $\bar{c}$ and value is $\underline{v}$. Therefore, in the fully efficient equilibrium, $q^{t_s,t_b}(\bar{c}, \underline{v}) = 0$, $q^{t_s,t_b}(\underline{c}, \underline{v}) = q^{t_s,t_b}(\underline{c}, \bar{v}) = q^{t_s,t_b}(\bar{c}, \bar{v}) = 1$ for all t_s and t_b.

2.2.1 Equilibria without Naive Traders

We consider first the setting without naive traders, i.e., $\varepsilon = 0$. The following proposition characterizes fully efficient equilibria in this case.

Proposition 1. *Suppose $\varepsilon = 0$. Strategies (Z_s^{str}, Z_b^{str}) are a fully efficient equilibrium if and only if $Z_s^{str}(\underline{c}) = Z_b^{str}(\underline{v}) = z_1$ and $Z_s^{str}(\bar{v}) = Z_b^{str}(\bar{c}) = z_2$, where*

$$\frac{2}{3}\underline{c} + \frac{1}{3}z_2 \le z_1 \le \underline{v} \qquad \textit{and} \qquad \bar{c} \le z_2 \le \frac{2}{3}\bar{v} + \frac{1}{3}z_1.$$

The proof is in Appendix A. The fully efficient equilibrium is achieved by the low-cost seller and the low-value buyer (and correspondingly the high-cost seller and the high-value buyer) agreeing on one price z_1 (and correspondingly z_2) to trade between them. Then there is also trade if the low-cost seller meets the high-value buyer and no trade when the low-value buyer meets the high-cost seller. The inequalities in the proposition ensure that the traders do not find it profitable to deviate from z_1 and z_2. Depending on the values $(\underline{c}, \underline{v}, \bar{c}, \bar{v})$, the fully efficient equilibrium may or may not exist.

Irrespective of whether a fully efficient equilibrium exists or not, there are multiple equilibria in this game, both in pure and mixed strategies. Any pure-strategy equilibrium falls into one of the following four categories: (a) a no-trade equilibrium, e.g., when both types of the buyer bid $\underline{c}$ and both types of sellers ask $\bar{v}$; (b) an equilibrium in which the high-value buyer trades with both types of the seller by coordinating on a price between $\bar{c}$ and $\bar{v}$ whereas the low-value buyer does not trade; (c) an equilibrium in which the low-cost seller trades with both types of the buyer by coordinating on a price between $\underline{c}$ and $\underline{v}$ whereas the high-cost seller does not trade; (d) an equilibrium in which the low-cost seller trades with the high-value buyer by coordinating on a price between $\underline{v}$ and $\bar{c}$ whereas the high-cost seller and the low-value buyer do not trade. Each of these four categories of pure-strategy equilibria is nonempty for all values of $(\underline{c}, \underline{v}, \bar{c}, \bar{v})$.

In a fully efficient equilibrium, the probability of trade for different costs and values is either 0 or 1. In any properly mixed equilibrium, the probability of trade will be strictly between 0 and 1 for some costs and values. Full efficiency cannot be achieved in a mixed equilibrium, and we thus focus only on pure equilibria of the double-auction game.

2.2.2 Equilibria with Naive Traders

Consider now the case with $\varepsilon \in (0,1)$. Define:

$$
\begin{aligned}
A(\underline{c}, \underline{v}) &= \frac{1}{3-\varepsilon}\left(2(1-\varepsilon)\underline{c} + (1+\varepsilon)\underline{v}\right)\\
B(z, \underline{c}, \underline{v}) &= \frac{1}{3-\varepsilon}\left(z + 2\underline{c} - \varepsilon\underline{v}\right)\\
C(z, \underline{c}, \underline{v}, \bar{v}) &= \frac{1}{3-\varepsilon}\left(\varepsilon\bar{v} + 2(2-\varepsilon)\underline{c} - \varepsilon\underline{v} - (1-\varepsilon)z\right)\\
D(\underline{c}, \underline{v}) &= \frac{1}{2-\varepsilon}\left(2(1-\varepsilon)\underline{v} + \varepsilon\underline{c}\right)\\
E(\bar{c}, \bar{v}) &= \frac{1}{2-\varepsilon}\left(2(1-\varepsilon)\bar{c} + \varepsilon\bar{v}\right)\\
F(\bar{c}, \bar{v}) &= \frac{1}{3-\varepsilon}\left(2(1-\varepsilon)\bar{v} + (1+\varepsilon)\bar{c}\right)\\
G(z, \bar{c}, \bar{v}) &= \frac{1}{3-\varepsilon}\left(z + 2\bar{v} - \varepsilon\bar{c}\right)\\
H(z, \bar{c}, \bar{v}, \underline{c}) &= \frac{1}{3-\varepsilon}\left(\varepsilon\underline{c} + 2(2-\varepsilon)\bar{v} - \varepsilon\bar{c} - (1-\varepsilon)z\right).
\end{aligned}
$$

The following proposition characterizes fully efficient equilibria when $\varepsilon \in (0,1)$.

Proposition 2. *Suppose $\varepsilon \in (0,1)$. Strategies (Z_s^{str}, Z_b^{str}) are a fully efficient equilibrium if and only if $Z_s^{str}(\underline{c}) = Z_b^{str}(\underline{v}) = z_1$ and $Z_s^{str}(\bar{c}) = Z_b^{str}(\bar{v}) = z_2$, where*

$$\max\{A(\underline{c}, \underline{v}), B(z_2, \underline{c}, \underline{v}), C(z_2, \underline{c}, \underline{v}, \bar{v})\} \leq z_1 \leq D(\underline{c}, \underline{v})$$

and

$$E(\bar{c}, \bar{v}) \leq z_2 \leq \min\{F(\bar{c}, \bar{v}), G(z_1, \bar{c}, \bar{v}), H(z_1, \bar{c}, \bar{v}, \underline{c})\}.$$

The proof is in Appendix A. As in the previous case without naive traders, a fully efficient equilibrium with naive traders is achieved by the low-cost strategic seller and the low-value strategic buyer (and correspondingly the high-cost strategic seller and the high-value strategic buyer) agreeing on one price z_1 (and correspondingly z_2) to trade between them. But now there are more deviations to consider, since instead of agreeing to trade at the price proposed by her strategic counterpart, a strategic trader may attempt to extract full surplus from her naive counterpart. The inequalities ensure that the strategic traders do not find it profitable to deviate from the strategies leading to fully efficient trade.

Since there are now more possible deviations, the conditions for the existence of fully efficient equilibria are more stringent that without naive traders. Even though a fully efficient equilibrium may exist without naive traders, it may cease to exist if a positive proportion of naive traders is present.

Similar to the case without naive traders, there are other equilibria, both pure and mixed. Since Proposition 2 is a characterization, those other equilibria do not achieve fully efficient trade in the double auction with naive traders.

2.3 Settings in the Experiment

We use the results from Propositions 1 and 2 to choose the parameter settings in the experiment. In the main treatment, we consider $\underline{c} = 10$, $\bar{c} = 70$, $\underline{v} = 30$, $\bar{v} = 90$. Without naive traders, there is a unique fully efficient equilibrium in this setting. If naive traders are present in any proportion, the fully efficient equilibrium disappears.

Proposition 3. *Suppose that $\underline{c} = 10$, $\bar{c} = 70$, $\underline{v} = 30$, $\bar{v} = 90$.*

i. If $\varepsilon = 0$, there exists a unique fully efficient equilibrium $Z_s^{str}(10) = Z_b^{str}(30) = 30$ and $Z_s^{str}(70) = Z_b^{str}(90) = 70$.

ii. If $\varepsilon \in (0, 1)$, there is no fully efficient equilibrium.

The proof is in Appendix A. It shows that without naive traders there is only one set of values z_1 and z_2 that satisfy the conditions for the existence of a fully efficient equilibrium given in Proposition 1. In this equilibrium, the low-value buyer and the high-cost seller obtain zero surplus, since they trade at price equal to their value or cost. The equilibrium is thus a knife-edge case in the sense that the low-value buyer and high-cost seller have only weak incentives to agree to such trades.

This knife-edge property of the fully efficient equilibrium in this case explains why full efficiency in not attainable in equilibrium in the presence of naive traders. Instead of agreeing to bid her value, the strategic low-value buyer can increase her surplus by shading her bid and trading with the naive low-cost seller without losing on trades with the strategic seller. Similarly, the strategic high-cost seller would like to exaggerate her ask to get more from the naive high-value buyer. The presence of naive traders thus makes these strategic traders behave more strategically (increase shading/exaggeration). As a result, there are no values for z_1 and z_2 that ensure that trade happens between all pairs of traders for whom there are positive gains from trade. Therefore, if the players manage to coordinate on the fully efficient equilibrium in the absence of naive traders, then the introduction of naive traders in the experiments should reduce efficiency.

Since there are multiple equilibria in the double auction (including those asymmetric between the seller and buyer roles), what equilibrium is played, if at all, may depend on many factors. Also, from a behavioral point of view, players may not necessarily coordinate on an equilibrium, act strategically or have selfish preferences. However, previous experiments on the double auction show that players coordinate on the most efficient equilibrium (in the setting with uniform distributions of values and costs on the same interval, see Radner and Schotter, 1989; Valley et al., 2002; McGinn et al., 2003; Ellingsen et al., 2009). We therefore expect to observe (close to) full efficiency in the setting without naive traders and the concomitant comparative static effect that the artificial introduction of naive traders reduces efficiency in the double auction.

Nevertheless, the knife-edge property of the fully efficient equilibrium in this setting could result in outcomes contrary to our expectation. First, the players might fail to coordinate on the fully efficient equilibrium when $\varepsilon = 0$ because, in this setting, some types of traders earn zero surplus in this equilibrium. Second, (some) players might expect (some) others to play naively, either because those others are not expected to understand strategic incentives of the game, or they are expected to care about honesty or efficiency. Thus, even before the artificial introduction of naive traders in the experiment, some players' underlying beliefs might be that a positive proportion of opponents are naive.

We therefore consider another setting ($\underline{c} = 5$, $\bar{c} = 70$, $\underline{v} = 30$, $\bar{v} = 95$), which tackles these two issues. The alternative setting has several fully efficient equilibria that result in positive surplus for all types of traders. At the same time, some of these fully efficient equilibria remain for small proportions of naive traders – thus making it possible to obtain full efficiency even when players' underlying beliefs are that a positive proportion of opponents are naive. But it is still the case that for sufficiently large presence of naive traders there is no fully efficient equilibrium, and thus the artificial introduction of naive traders, as we do in the experiment, may reduce efficiency even in this alternative setting.

Proposition 4. *Suppose that* $\underline{c} = 5$*,* $\bar{c} = 70$*,* $\underline{v} = 30$*,* $\bar{v} = 95$*.*

i. If $\varepsilon = 0$*, there exist several fully efficient equilibria.*

ii. If $\varepsilon \in (0, 1)$*, a fully efficient equilibrium exists for* $\varepsilon \in (0, \bar{\varepsilon}]$*, where* $\bar{\varepsilon} = (11 - \sqrt{101})/5 \approx 0.190$*, and does not exist for* $\varepsilon \in (\bar{\varepsilon}, 1)$*.*

The proof is in Appendix A. Fully efficient equilibria in this case exist for all values of ε

up to approximately 0.19. When ε exceeds 0.19, the strategic traders' incentives to extract surplus from the naive traders becomes overwhelming, and it is not possible to achieve efficient trade between the strategic traders anymore.

All the above propositions are based on self-interested risk-neutral strategic traders. However, many plausible alternative formulations of preferences preserve a fully efficient equilibrium whenever it exists. For example, if players are risk-averse, a fully efficient equilibrium remains, since any deviation from it either reduces a player's ex-post payoff or increases the variance of the expected payoff. Similarly, if players are efficiency concerned, the fully efficient equilibrium remains. If players have a cost of lying (setting price different from their value/cost), a fully efficient equilibrium remains if the cost is not large.[11]

3 Experiment Design

We consider experimental treatments differing along the following dimensions: the absence of presence of naive traders, different parameter settings discussed above, and the framing of interactions. Previous experimental results on double auctions indicate that (second-best) efficiency is achievable even without communication (in the setting with a uniform distribution of values and costs on the same interval, see Radner and Schotter, 1989; Valley et al., 2002; McGinn et al., 2003; Ellingsen et al., 2009). We therefore focus on the double auction without communication, although we also run some sessions with communication.

Players in the double auction make decisions knowing the distribution of possible costs and values, and knowing their own cost (if seller) or value (if buyer). Our main treatment difference is in the artificially induced proportion of naive traders. In treatments labelled "S" (for "strategic") all subjects are matched in pairs and play the double auction. In treatments "N" (for "naive"), subjects are told that with probability 0.25 their decision will be matched with that of a computerized opponent who sets bid equal to the buyer's value or ask equal to the seller's cost. This arrangement corresponds to the case with proportion $\varepsilon = 0.25$ of naive traders.[12]

The distributions of costs and valuations follow the two settings discussed in the previous section. One setting involves $\underline{c} = 10$, $\underline{v} = 30$, $\bar{c} = 70$, $\bar{v} = 90$ (coded as "20" by the difference

[11] All these alternative preferences can create additional fully efficient equilibria, also for the setting with naive traders. However, as we will see below, fully efficient equilibria are not what is observed in the experiment.

[12] To be precise, subjects are still matched in pairs in this treatment. After they have chosen their bid/ask, with probability 0.25, both subjects in a pair are rematched with computerized opponents who play naively. Thus, each subject in the pair faces a 0.75 probability of being matched with the human opponent and 0.25 probability of being matched with a naive opponent. Hence, the strategic incentives in the "N" treatment correspond to the case with $\varepsilon = 0.25$ of naive traders.

$\underline{v} - \underline{c} = \bar{v} - \bar{c} = 20$.) Recall that in this setting there exists a fully efficient equilibrium if the proportion ε of naive traders is 0, but not for any other value of $\varepsilon < 1$. Increasing ε by 0.25 as is done in the "N" treatment may thus decrease efficiency compared to the "S" treatment. However, as mentioned earlier, we might fail to observe the fully efficient equilibrium in this setting for two reasons: firstly, because some types of traders earn zero surplus (making it potentially difficult to coordinate on this equilibrium); secondly, it is possible that (some) subjects' underlying beliefs about opponent's naiveté are positive. We therefore also consider setting with $\underline{c} = 5$, $\underline{v} = 30$, $\bar{c} = 70$, $\bar{v} = 95$ (coded as "25" for $\underline{v} - \underline{c} = \bar{v} - \bar{c} = 25$). In this setting, for values of ε less than 0.19, there exist fully efficient equilibria in which all types of traders earn positive surplus. Thus, the "S" treatment in this setting offers the opportunity to attain full efficiency with positive surplus to all types of traders as long as the subjects' underlying beliefs regarding their opponent's naiveté are not too high. But full efficiency becomes unattainable in equilibrium once the proportion of naive traders is artificially increased by 0.25 in the "N" treatment. Our hypothesis is that efficiency in treatment "S" is higher than in treatment "N" because (at least some) human traders behave more strategically with a higher proportion of naive traders.

We also vary the framing of the interaction. In the double-auction framing (coded "BA"), the seller names an ask price and the buyer names a bid price. In addition to this standard framing, we consider the direct mechanism framing (coded "DM"). In this framing, subjects are asked to report their cost or value. The payoffs are calculated according to the double-auction rule: trade occurs if the reported value is above the reported cost, at the price in the middle between the reported value and cost. The difference between BA and DM frames is only in the wording of the instructions (a sample of which is given in Appendix B). Framing the interaction as a direct mechanism makes it clearer what revealing one's cost or value is. Our expectation is that experiment participants in frame DM play naively more often than in frame BA. If other participants realize this, it can have a knock-on effect on their behavior, making them act more strategically in frame DM than in frame BA, and thus efficiency may be higher in frame BA than in frame DM.

Each experimental session without communication follows one of the settings. For example, a session coded as "BA-S-20" means a session using the double-auction frame, no artificial naive players, and parameters $\underline{c} = 10$, $\underline{v} = 30$, $\bar{c} = 70$, $\bar{v} = 90$. In each session, 16-18 subjects are divided into two matching groups of 8-10 players. Within each matching group, in each period half of the subjects have one role (e.g. Seller) and the other half the other role. Table 1 gives the number of sessions, matching groups, and subjects for each treatment. In total, there are 16 sessions with 272 subjects.

A session lasts 40 periods. In each period, a subject in one role is randomly matched

Code	Sessions	Matching groups	Subjects	Code	Sessions	Matching groups	Subjects
BA-S-20	3	6	52	DM-S-20	2	4	34
BA-N-20	3	6	52	DM-N-20	2	4	32
BA-S-25	2	4	34	DM-S-25	1	2	18
BA-N-25	2	4	32	DM-N-25	1	2	18

Table 1: Experimental sessions without communication

with a subject in the other role. During a session, each subject keeps the same role for 10 periods, then switches to the other role.[13] After the matching, each subject's value or cost is drawn randomly and independently. The subjects then make their decisions[14] and their payoffs are calculated.[15] At the end of each period, subjects are told what the outcome in their match is (but neither the opponent's cost or value nor whether the opponent is a computer). Payoffs are measured in points.

Experimental sessions were conducted in the Centre for Decision Research and Experimental Economics (CeDEx) laboratory at the University of Nottingham. The experiment was programmed in z-Tree (Fischbacher, 2007) and subjects were recruited from the CeDEx database of experimental participants using the ORSEE software (Greiner, 2015). They were students of various disciplines across the university. Together with reading the instructions, answering control questions, and filling in the post-experiment questionnaire, each 40-period session lasted 90-120 minutes. Subjects were paid their accumulated earnings, converted from points to pounds at the rate £0.20 for 10 points in setting "20" and £0.15 for 10 points in setting "25" (to equalize payoffs, since the available surplus is larger in setting "25"). The average payment per subject was £14.12 ($22.25 at the time of the experiment), including a £5 show-up fee.

The previous experiments on double auction that also considered pre-play face-to-face communication (e.g. Valley et al., 2002 and McGinn et al., 2003) found that it increases efficiency beyond the second-best. We therefore also run two sessions of treatment "BA-S-20" with face-to-face communication. There are 16 subjects in each session, 8 sellers and 8 buyers in separate rooms. After their costs and values are determined, subjects are taken

[13]Thus a subject could be e.g. Seller in periods 1-10, Buyer in periods 11-20, Seller again in periods 21-30 and Buyer again in periods 31-40.

[14]Before making a decision, subjects could use a "Payoff Calculator" to calculate their own expected payoffs for different possible asks/bids set by themselves and their opponents.

[15]In the N treatment, it is first determined whether the subjects in a pair are rematched against a computerized opponent and then payoffs are calculated.

to another room to have a 3-minute face-to-face conversation. Then they return to their rooms and make decisions about bid/ask price. Subjects are matched in a round-robin format during 8 periods and never meet again a subject with whom they already met face-to-face. Subjects were paid £1 for 10 points in these treatments. These sessions lasted 150 minutes and subjects earned on average £19.25 ($30.13; the payoff is higher partially because the random draws favoured low costs and high values). We expect these sessions to achieve at least as high (close to full) efficiency as the sessions without communication.

4 Experiment Results

4.1 Efficiency

Before we present our results, we define how we measure efficiency in our treatments. The *realized efficiency* is the proportion of total available surplus that is captured by the traders in our game. In treatment S, the realized efficiency is the ratio of the total realized gains from trade to the total available gains from trade. However, the realized efficiency in treatment N is muddled by the outcome of the matching lottery (whether the subjects are matched with computerized opponents or not) and hence is not directly comparable to the realized efficiency in treatment S which only depends on the realization of costs and values and bids and asks. To make efficiency measures comparable, we construct two measures for treatment N using only the realized values and costs and the subjects' bids and asks.

The first measure of efficiency in treatment N corresponds to the proportion of total available surplus that is expected to be captured by the traders after the strategic traders have chosen their bids and asks but before the disposition types of the buyer and seller are known. We thus calculate the expected gains from trade in each matching pair given the realized cost and value and the subjects' bid and ask.[16] Then the *expected efficiency* in treatment N is the ratio of the total expected gains from trade to the total available gains.

The second measure of efficiency in treatment N corresponds to the proportion of total available surplus that would have been captured by the traders if hypothetically all subjects were matched between themselves rather than with (possibly) a computerized opponent. Thus, the *human efficiency* in treatment N is the ratio of the total gains from trade that would have realized had the subjects in each pair matched to each other, to the total available gains from trade.

If applied to treatment S, both the expected efficiency measure and the human efficiency

[16] For example, if the buyer's value is 90 and the seller's cost is 10, and the subject in the buyer's role bids 50 while the subject in the seller's role asks 75, then the expected gains from trade are $(0.75)^2(0) + 2(0.75)(0.25)(80) + (0.25)^2(80) = 35$.

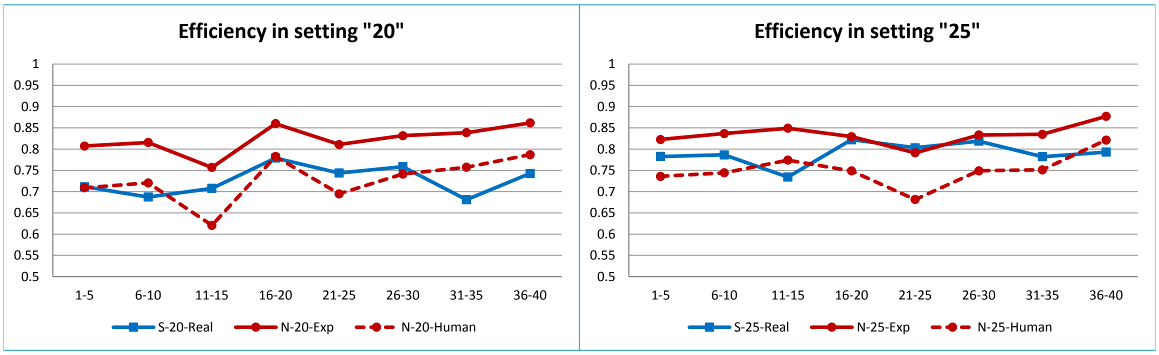

Figure 1: Efficiency comparison in treatments S and N

measure coincide, and they are also equal to the realized efficiency measure. The theoretical analysis of the situations played out in the experiment shows that if the strategic traders coordinate on a fully efficient equilibrium in the absence of naive traders, then efficiency (both expected and human) can be higher without naive traders (treatment S, $\varepsilon = 0$) than with a sizable proportion of naive traders (treatment N, $\varepsilon = 0.25$). Also, if there are more naive plays in frame DM than in frame BA, then it is possible that efficiency is higher in the BA frame than in the DM frame.

Figure 1 presents the comparison of efficiency in treatments S and N. It shows the time series of the three efficiency measures – realized efficiency in treatment S labeled "S-Real" (blue, lighter line), expected efficiency in treatment N labeled "N-Exp" (red, darker line), and human efficiency in treatment N labeled "N-Human" (dotted line) – averaged over blocks of five periods. The efficiency measures are combined across the BA and DM frames, since none of the non-parametric Mann-Wilcoxon-Whitney rank-sum tests on the level of matching groups comparing median efficiency shows significant difference across frames.[17]

As can be seen in the figure, all measures of efficiency are clearly lower than full efficiency and do not change much over time. The expected efficiency in treatment N appears higher than the realized efficiency in treatment S, especially in setting "20". Non-parametric tests (focusing on periods 21-40, allowing subjects to learn the game) find a significant difference between these measures of efficiency in the S and N treatments in setting "20" (p-value of the two-sided test 0.023; 10 observations for each treatment) but not in setting "25" (p-value 0.337; 6 observations for each treatment). There appears to be little difference between the treatments when comparing the realized efficiency in treatment S with the human efficiency in treatment N. Indeed, non-parametric tests do not find significant differences for either of

[17]The tests results are presented in Section C.1.1 of Appendix C.

Gains	Setting “20” S	Setting “20” N (N-Human)	Setting “25” S	Setting “25” N (N-Human)
$\bar{v} - \underline{c}$	0.853	0.922 (0.861)	0.947	0.947 (0.906)
$\bar{v} - \bar{c}$	0.551	0.676 (0.528)	0.622	0.650 (0.493)
$\underline{v} - \underline{c}$	0.436	0.613 (0.438)	0.479	0.646 (0.479)
0 ($\bar{c} > \underline{v}$)	0.015	0.018 (0.023)	0.008	0.043 (0.050)

Table 2: Probabilities of trade for different available gains

the comparisons (p-values are 0.650 for parameter setting “20” and 0.749 for “25”).

Result 1. *Efficiency*

- *Full efficiency is not realized even when it is possible in equilibrium.*
- *Expected efficiency is higher in treatment N than realized efficiency in treatment S, significantly so in setting “20”.*
- *Human efficiency in treatment N is not significantly different from realized efficiency in treatment S.*
- *There are no significant differences in efficiency between BA and DM frames.*

Thus our hypotheses related to efficiency are not confirmed. Indeed, the results go in the opposite direction: the expected efficiency in treatment N is higher than the realized efficiency in treatment S. Since the human efficiency in treatment N is not different from the realized efficiency in treatment S, the greater value of the expected efficiency in treatment N compared to the realized efficiency in treatment S is mostly due to additional trades with and between computerized naive traders rather than between human subjects.

The efficiency measures presented in Figure 1 combine situations with a high available surplus ($\bar{v} - \underline{c}$) and a smaller surplus ($\underline{v} - \underline{c}$ and $\bar{v} - \bar{c}$) as well as situations with zero surplus ($\bar{c}$ and $\underline{v}$). Table 2 shows the probabilities of trade for these different situations averaged across all rounds (combined across BA and DM frames). In the S treatment, we measure the realized probability of trade given the subjects' bids and asks (labeled “S”). In the N treatment, we again distinguish between the expected probability of trade given the subjects' bids and asks and the realized costs and values (labeled “N”) and the human probability of trade if all subjects were always matched between themselves rather than (occasionally) with a computerized opponent (labeled “N-Human”).

From Table 2, the expected probability of trade is higher in treatment N than the realized probability of trade in treatment S for all gains from trade (even for the negative ones $\underline{v} - \bar{c}$, although such trades happen only in less than 3% of all cases of negative surplus).[18] However, the human probability of trade in treatment N is very similar to the realized probability of trade in treatment S for all gains of trade.[19] The introduction of naive traders thus raises the expected probability of trade in all situations with positive surplus but does not change the probability of trade in matches between human subjects.

An obvious reason for the failure to get a higher efficiency in the S treatment compared to the N treatment is that subjects are not playing a fully efficient equilibrium. The absence of a significant difference between efficiency in the S treatment and efficiency in matches of human subjects in the N treatment, and between the probabilities of trade in treatment S and in matches of human subjects in treatment N, points also to the possibility that subjects do not behave more strategically in treatment N compared with treatment S. This explanation is discussed in more detail in the next section.

In the communication treatment (which uses parameter setting "20"), the realized efficiency is 0.884, higher than in the treatments without communication but still falling somewhat short of full efficiency.[20] The realized probability of trade for the large surplus $\bar{v} - \underline{c}$ is 0.939, for surplus $\bar{v} - \bar{c}$ it is 0.788 and for $\underline{v} - \underline{c}$ it is 0.621 (and 0 for negative surplus). Even though trade almost always happens in the case of large surplus, subjects do not always manage to trade in the cases of lower surplus, leaving some surplus unrealized.

4.2 Individual Behavior

We divide the analysis of the individual choices into two parts. First we look how often the choices coincide with cost (for sellers) and value (for buyers), to estimate the proportion of naive play in the experiment. In the second part we look at strategic choices (i.e., ones that are different from cost or value) to determine the extent of shading or exaggeration.

4.2.1 Incidence of Naive Behavior

We term as "naive" behavior that corresponds to setting bid equal to value as a buyer and ask equal to cost as a seller (in the DM frame, this corresponds to reporting the actual value or cost). Naive behavior is weakly dominated from purely selfish preferences point of view: reducing the bid (or raising the ask) by the smallest available unit (0.01 in the experiment) does not lose any trades with a positive surplus for the player while leading

[18] It is also the case that the probability of trade with $\bar{v} - \bar{c}$ is larger than with $\underline{v} - \underline{c}$. The explanation for this is discussed in the next subsection.

[19] These observations are confirmed by regressions reported in Section C.1.2 of Appendix C. The regressions also find that the probabilities of trade are not affected by frame (BA or DM) and setting ("20" or "25").

[20] A summary of the data for the communication sessions is provided in Section C.3 of Appendix C.

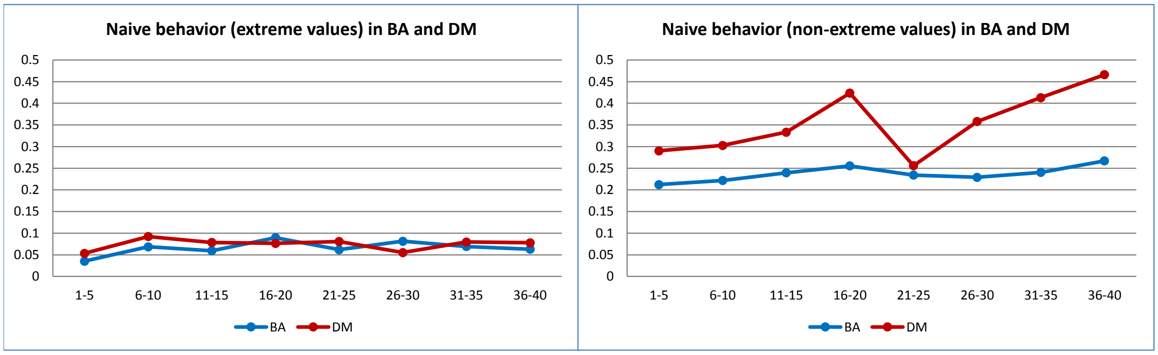

Figure 2: Incidence of naive behavior

to a better price. However, in the S treatment, such behavior by the low-value buyer and high-cost seller ("non-extreme" values) is consistent with a fully efficient equilibrium, and thus not necessarily irrational. But in the N treatment, such behavior is not consistent with any equilibrium since the low-value buyer (high-cost seller) can obtain a positive expected payoff by bidding equal to $\underline{c}$ (asking equal to $\bar{v}$). We can therefore expect that the incidence of naive behavior is higher in treatment S than in treatment N, at least for traders with non-extreme values. We can also expect the DM frame to increase the incidence of naive behavior due to, e.g., aversion to lying.

Figure 2 plots the incidence of naive behavior over time (in blocks of 5 periods), disaggregated by frames BA and DM (but aggregated across other treatment differences)[21]. The left panel is for the extreme values and the right panel is for the non-extreme values.

For the extreme values, there is almost no change in the proportion of honest play, and there is no difference between the BA and DM frames. The overall frequency of naive behavior with extreme values stays around 7% throughout the experiment. If this proportion is taken as an indicator of naturally occuring naive play, it allows the existence of a fully efficient equilibrium in setting "25".

For the non-extreme values, the figure shows that the proportion of naive play is higher in frame DM than in frame BA. Non-parametric tests find that the difference is significant for periods 21-40 (p-value of the two-sided Mann-Wilcoxon-Whitney test is 0.016), partially confirming our hypothesis. The proportion of naive play also appears to increase in the DM frame over time.

The aggregate proportions presented in the figure may hide other possible differences.

[21] Non-parametric tests comparing the proportions, reported in Section C.2.1 of Appendix C, do not find significant differences across those treatment differences.

Determinants of naive behavior in Periods 21-40 Dependent variable: $NaivePlay$		
Variable	Extreme	Non-extreme
$NPresent$	0.023** (0.012)	0.032 (0.032)
DM	−0.001 (0.012)	0.095*** (0.030)
$Val25$	0.009 (0.012)	0.028 (0.030)
$Seller$	−0.036*** (0.010)	0.063*** (0.017)
$Period$	< 0.001 (0.001)	0.004*** (0.001)
$NaivePlay_{t-1}$	0.134*** (0.024)	0.278*** (0.034)
$PayBRHist$	0.001 (0.002)	−0.012 (0.014)
$PayBRLast$	< 0.001 (< 0.001)	−0.007*** (0.002)
($Constant$)	−2.343*** (0.401)	−2.221*** (0.260)
Observations	2708	2732
Clusters	272	272

Average marginal effects of variables presented. Standard errors clustered by subjects in parentheses. *** - significant at 1% level; ** - significant at 5% level.

Table 3: Determinants of naive behavior

We try to get a further insight into the determinants of naive (and strategic, in the next subsection) behavior by looking at individual decisions. To this end, we regress subjects' choices on variables describing the treatment and additional variables attempting to capture a subject's experienced history in the game.

Specifically, we consider the dummy variable $NaivePlay_{it}$, equal to 1 if the play of subject i at period t is naive. Among the explanatory variables, $NPresent_i$, DM_i, $Val25_i$ are the dummy variables describing the treatment, frame, and setting, and $Seller_{it}$ is the dummy variable equal to 1 if the subject is a seller in period t.

The history variables are constructed as follows. One is the lagged naive play dummy, $NaivePlay_{it-1}$. For each subject i, variables $PayBRHist_{it}$ and $PayBRLast_{it}$ measure the payoff difference between playing the best response (in the given role) and playing naively. Variable $PayBRHist_{it}$ measures it for the whole history of play up to period t, while variable $PayBRLast_{it}$ measures it against the last match of subject i. The variables intend to capture incentives to deviate from naive play, both for a long and a short memory.

We run random-effects probit regressions, separately for the subsamples of observations with the extreme and with the non-extreme values. The results, for periods 21-40, are presented in Table 3.

The regressions confirm the significantly higher probability to observe naive plays in frame DM than in frame BA as well as the upward trend in this proportion for traders with non-extreme values. They also show that the presence of naive traders leads to a significantly higher probability of naive play for the extreme-value traders but not for the non-extreme value traders.[22] We also find that with the extreme values, sellers play naively significantly less often than buyers (the raw proportions are 4% versus 10%). With the non-extreme values the result is reversed: buyers play naively significantly less often than sellers (25% versus 34%).[23]

The incentives to deviate from naive play, as measured by the history variables, do not seem to be important for extreme value traders although one of them, $PayBRLast$, plays a role in the decisions of non-extreme value traders. Thus, a higher achievable payoff from best response to the last observation decreases the probability of naive play for the non-extreme value traders. The largest effect, though, is from the previous decisions to play naively, and determines such later decisions for both extreme and non-extreme value traders, suggesting that naivety may be a consistent behavioral trait. However, we cannot clearly demarcate subjects into strategic and naive types. While 24% of subjects were always strategic, none of the subjects always played naively.

One way to explain both the greater frequency of naive play among the non-extreme value traders and the different impact of an increase in the payoff from best responding on non-extreme-value versus extreme-value traders is that subjects have (non-monetary) costs of lying. To see this in the simplest possible way, consider an alternative model in which all players are strategic but the buyer pays a cost of x if she lies.[24] Suppose the low-cost seller asks z_1 with $\underline{c} \leq z_1 < \underline{v}$ while the high-cost seller asks z_2 with $\bar{c} \leq z_2 < \bar{v}$. Then the low-value buyer gains $\frac{1}{4}(\underline{v} - z_1) - x$ by lying (setting bid equal z_1) whereas the high-value buyer gains $\frac{1}{4}(z_2 - z_1) - x$ by such a lie. As the latter is greater than the former, we can thus have a situation where the low-value buyer is acting naively whereas the high-value buyer is lying – which is consistent with a greater proportion of naive play by the traders with non-extreme values in the experiment. A decrease in z_1, which increases the payoff from lying for both types of buyers, may cause the low-value buyer to switch to lying – which is consistent with the reduction in naive behavior by the traders with non-extreme values in response to an increase in the payoff from best responding.

That the non-extreme value traders have a lower expected gain from lying than the

[22]The absolute proportions of extreme-value traders playing naively is small though: 5% in treatment S and 8% in treatment N.

[23]Radner and Schotter (1989) also find a difference in buyer and seller behavior in their experiment on symmetric double auction.

[24]For example, Ellingsen and Johannesson (2004) consider a model with such a fixed cost of lying.

extreme-value traders could similarly explain more naive play by the former but not by the latter types of traders in the DM frame if framing the actions as a report about one's cost/value increases the cost of lying. Having said that, for cost of lying to explain the differences in the observed naive behavior across traders' roles would require us to assume that the extreme-value sellers have a lower cost of lying than the extreme-value buyers while the non-extreme-value sellers have a higher cost of lying than the non-extreme-value buyers. It is a priori not clear why that should be the case since the experimental setting is symmetric across traders' roles.

Summarizing these observations, we have:

Result 2. *Naive behavior*

- *Subjects are more likely to play naively if they played naively in the past.*
- *With extreme values, (a) naive play is observed in 7% of decisions and there is little difference in the frequency of naive play across treatments or frames; (b) buyers are more likely to play naively than sellers.*
- *With non-extreme values, (a) there is more naive play in frame DM than in frame BA but there is no difference between treatments S and N; (b) sellers are more likely to play naively than buyers.*

Our expectation to see more naive plays in treatment S than in treatment N is not confirmed. On the other hand, the expectation that the DM frame is more conducive to naive play than the BA frame, is partially confirmed (for the non-extreme but not for the extreme-value traders). More surprising is the difference across traders' roles since by design all subjects played an equal number of times in each of the roles and were equally likely to receive extreme or non-extreme valuations. That the high-value buyers and the high-cost sellers play naively more often than the low-value buyers and the low-cost sellers provides an explanation of the difference in probabilities of trade in cases $\bar{v} - \bar{c}$ and $\underline{v} - \underline{c}$ that was noted in the previous subsection.

In the communication session, naive behavior in the double-auction stage is much less common: 3% with the extreme values and 9% with the non-extreme values. However, transcripts of subjects' conversations show that true values and costs are often revealed during the communication stage. Almost all traders with non-extreme values reveal them truthfully (94%). Of the traders with extreme values, 56% do so.[25] After exchanging value/cost statements (which happens in 91% of conversations, out of which 59% involve

[25] Of the remaining extreme-value traders, 36% lied about their value/cost misstating it as non-extreme, and 8% did not state it unambigously. This can be seen in Section C.3 of Appendix C.

mutual truthful revelation), subjects often agree on one price (in 89% of cases with the stated value and cost allowing positive gains from trade), usually in the middle of their stated values/costs (96% of cases). These one-price agreements are usually mutually followed (86% of cases). Therefore, "naivety" has a different meaning in the communication treatment: it is whether to reveal true cost or value in the communication stage rather than set ask or bid equal to cost or value in the price-setting stage.

4.2.2 Extent of strategic behavior

The previous section looked at the incidence of naive behavior, which also explains the incidence of strategic behavior, i.e., how often subjects bid or ask anything other than their value or cost. In this section, we analyze the extent of strategic behavior, which measures how far bids and asks are from values or costs, i.e. we study the extent of shading (if buyer) or exaggeration (if seller) in the experiment.

Compared with a fully efficient equilibrium in treatment S, the presence of naive traders in treatment N makes strategic traders with non-extreme values increase shading of bids or exaggeration of costs, since they can get a higher expected payoff by exploiting naive traders. Our hypothesis is thus to observe a higher extent of strategic behavior by the traders with non-extreme values in treatment N than in treatment S (and possibly in frame DM, where more naive traders can be present, than in frame BA).

The effect of introduction of naive traders on the equilibrium behavior of strategic traders with extreme values is less clear cut. For example, in the unique fully efficient equilibrium of the S-20 treatment the high-value buyer and the high-cost seller coordinate on the price of 70 while the low-value buyer and the low-cost seller coordinate on the price of 30. In the N-20 treatment, where $\varepsilon = 0.25$, there exists an equilibrium in which the high-value buyer and the low-cost seller coordinate on the price of 28, the low-value buyer bids 10 and the high-cost seller asks 90. Compared to the fully efficient equilibrium in the S-20 treatment, the low-cost seller now acts less strategically by reducing her ask whereas all other types of traders act more strategically by increasing the amount of shading or exaggeration.[26] Due to this ambiguous effect, we do not state a formal hypothesis comparing the extent of strategic behavior by the traders with extreme values.

Figure 3 shows the time series of the average amount of shading and exaggeration (conditional on asks not equal cost and bids not equal value). The data are separated for

[26]There is also an equilibrium in which the high-value buyer and the low-cost seller coordinate on the price of 72, the low-value buyer bids 10 and the high-cost seller asks 90. Here, compared to the fully efficient equilibrium in the S-20 treatment, the high-value buyer acts less strategically by increasing her bid whereas all other types of traders act more strategically.

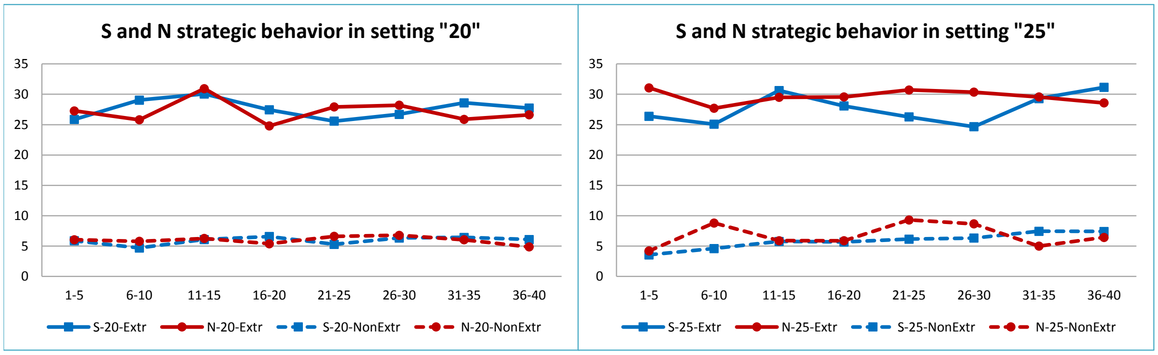

Figure 3: Extent of strategic behavior

the "20" and "25" settings, and disaggregated between S (labeled 'S-') and N (labeled 'N-') treatments (while aggregated across BA and DM frames)[27]. The solid lines are for the extreme-value traders while the dotted lines are for non-extreme-value traders.

Unsurprisingly, traders with the extreme values shade or exaggerate much more than traders with the non-extreme values. From the figure, there appears to be little difference in the average extent of strategic behavior between treatments S and N, which is also confirmed by non-parametric tests (p-values of the comparisons are 0.762 for the extreme values and 0.880 for the non-extreme values in setting "20" and 0.873 for the extreme values and 0.423 for the non-extreme values in setting "25"). The extent of strategic behavior also stays roughly at the same level across periods. Thus our expectation of more strategic behavior in treatment N than in treatment S (or in frame DM than in frame BA) is not confirmed.

To get a further insight into determinants of strategic behavior from individual data, we regress the variable measuring the extent of strategic behavior on variables describing the treatment setting and a subject's history. Strategic behavior (variable $StratBeh$) is measured as $Ask - Cost$ for sellers and $Value - Bid$ for buyers. The history variables are now summarized in the distance of best response (to the whole history of observations, variable $DistBRHist$, and to the last observation, variable $DistBRLast$) from the cost or value, since these variables measure the "optimal" strategic behavior in view of (long or short memory of) the player's history. The results of the random-effects linear regressions, again for separate subsamples of players with the extreme and the non-extreme values are presented in the first two columns of Table 4.

As with naive behavior, one of the main determinants of the extent of strategic behavior

[27]Non-parametric tests reported in Section C.2.2 of Appendix C do not find significant differences in strategic behavior between frames.

Determinants of the extent of strategic behavior in Periods 21-40				
	Dependent variable: $StratBeh$		Dependent variable: $\Delta StratBeh$	
Variable	Extreme	Non-extreme	Extreme	Non-extreme
$NPresent$	-0.117 (0.790)	-0.358 (0.838)	-1.242^{**} (0.567)	-2.156^{***} (0.428)
DM	0.538 (0.812)	-0.534 (0.900)	0.176 (0.584)	-0.758^{*} (0.407)
$Val25$	0.687 (0.841)	0.183 (0.886)	-0.458 (0.603)	-0.853^{**} (0.398)
$Seller$	0.583 (0.581)	-0.238 (0.347)	-0.112 (0.457)	-0.546^{**} (0.251)
$Period$	0.013 (0.050)	-0.062^{*} (0.032)	0.017 (0.041)	-0.051^{**} (0.022)
$StratBeh_{t-1}$	0.658^{***} (0.027)	0.188^{***} (0.051)		
$DistBRHist$	0.037^{**} (0.016)	0.116^{***} (0.044)	0.123^{***} (0.016)	0.256^{***} (0.035)
$DistBRLast$	0.002 (0.013)	0.025 (0.022)	0.066^{***} (0.011)	0.185^{***} (0.025)
$Constant$	7.480^{***} (1.877)	5.911^{***} (1.259)	-1.751 (1.380)	1.556^{**} (0.711)
Observations	2517	1931	2708	2731
Subjects	271	256	272	272

Standard errors clustered by subjects in parentheses. *** - significant at 1% level; ** - significant at 5% level; * - significant at 10% level

Table 4: Determinants of the extent of strategic behavior

of a player is how strategic the player was in the past: the more strategic the player was before, the greater the extent of current strategic behavior by the player (here, $StratBeh_{t-1}$ refers to the previous time the player was in a given role with the given cost/value). An increase in the distance between the best response to the entire history of play and the players' value/cost increases the extent of strategic behavior too, although only by a small proportion of the increase in the distance (and the immediate history does not have a significant effect). Variables describing treatment, frame and setting have no significant effect, confirming non-parametric tests.

The regressions for the level of strategic behavior do not clearly indicate how players change their choices because this change depends on the interplay of the coefficients on $StratBeh_{t-1}$ and $DistBRHist$. We thus also analyze what determines these changes. For this, we use variable $\Delta StratBeh = StratBeh_{it} - StratBeh_{it-1}$. The history variables are modified in a similar way (i.e., to $DistBRHist_{it} - StratBeh_{it-1}$ and $DistBRLast_{it} - StratBeh_{it-1}$), thus measuring how much strategic behavior should be changed if the subject is best responding to the observed history of asks or bids (or to the last observation). Both these variables try to capture incentives to increase or decrease shading or exaggeration.

The results of the regressions for changes in strategic behavior are presented in the last two columns of Table 4 (the regressions are simple linear regressions; all observations are included, also the naive ones since they represent zero strategic behavior which can change from one period to the next). The regressions show again that the history is important for changes in the extent of strategic behavior. Shading or exaggeration is larger if the best response to the observed history is further away from the current choice, also for the immediate history (last experienced observation). It also goes on average in the right direction: the mean value of $\Delta StratBeh$ is positive for positive values of $DistBRHist$ and $DistBRLast$ and negative for negative value of these history variables. Thus players react to incentives from the observed history (albeit only partially).

The regressions also find that the treatment variables have some effect on the magnitude of average changes in the extent of strategic behavior, especially for traders with the non-extreme values. The average levels of strategic behavior are similar in all treatments though, and stay constant over time. Summarizing,

Result 3. *Extent of strategic behavior*

- *The extent of strategic behavior in the past positively effects the extent of strategic behavior in future.*
- *There are no significant differences in the extent of strategic behavior across treatments and frames.*
- *Strategic behavior changes towards best responses to previous observations (both distant in the past and immediate) but only partially.*

Thus, subjects do not appear to behave more strategically in treatments with more naive players (treatment N as compared with treatment S; recall also that there were significantly more naive plays in the DM frame than in the BA frame). However, the choices do react to the previous experience in the manner one can expect, towards best responses to observations. We use these results in the next section to motivate a dynamic model with noisy best response that can account for (some of) the properties of the data from the experiment.

In the communication sessions (recall that they use parameter setting "20"), the extent of strategic behavior in the double-auction stage appears larger than in the sessions without communication: with the extreme values, the average strategic behavior is 30.5 and with the non-extreme values it is 11.7. This is, however, misleading, since traders often reveal their value/cost and agree on bid/ask in the middle between the revealed value and cost, as shown before. Indeed, the average strategic behavior of extreme-value traders is 37.7 if the other trader also has an extreme value,[28] but only 18.7 if the other trader has a non-extreme

[28] In such pairs, buyers's strategic behavior (average 47.7) is larger than sellers' (27.4). Buyers also lie about the value more in such pairs than sellers do, as can be seen in Section C.3 of Appendix C.

value.[29] Note that in the pairs in which an extreme-value trader meets a non-extreme value trader, the strategic behavior of the former is larger and not always compatible with that of the latter: since values and costs are sometimes misstated, the claims of non-extreme values are not always believed, leading to the efficiency losses discussed earlier.

5 A Dynamic Adjustment Model of Choices

A few points learned from the previous section about the behavior of the participants in the sessions without communication are that the average behavior is similar across settings and hardly changes over time, even though some reasonable adjustments in view of experience are made. In this section we present a dynamic adjustment model that replicates these features of the experimental data.

As the starting point of the model, we observe that in the first round the extreme-value traders exaggerate their cost or shade their value by 27.77 on average; the non-extreme-value traders do so by 4.35.[30] It turns out that if traders make choices based on low levels of strategic sophistication like in level-k model (Stahl and Wilson, 1994), the average play roughly fits these numbers.

More precisely, we consider the uniform strategy of one trader, discretized on integer values between 0 and 100 (the range used in the experiment). Suppose that the choices of the other trader follow a logistic distribution based on the expected payoffs against this uniform strategy, i.e.,

$$Prob(j) = \frac{e^{\lambda Eu(j)}}{\sum_{j=0}^{100} e^{\lambda Eu(i)}}, \tag{1}$$

where $j \in \{0, \ldots, 100\}$ and $Eu(i)$ is the expected payoff of strategy i against the uniform strategy of the other trader. Parameter λ measures the noise in choosing the strategy with the highest payoff: as $\lambda \to \infty$, the best response is chosen with probability approaching 1, while for $\lambda = 0$ all strategies are chosen with the same probability. For intermediate values of λ, strategies with higher payoffs have a higher probability to be chosen than strategies with lower payoffs.[31]

[29]For the non-extreme value traders, average strategic behavior is 9.9 if the other trader has an extreme value and 15.4 if the other trader has a non-extreme value. This difference is not as telling though since if value 30 and cost 70 are stated, any (individually rational) bids and asks are possible.

[30]The numbers are also similar if all rounds are considered: 25.97 for the extreme-value traders and 4.35 for the non-extreme-value traders. These numbers are slightly lower that those that can be seen in Figure 3, where observations with ask equal to cost and bid equal to value are excluded.

[31]This logistic choice is used in the level-k model of choices (Stahl and Wilson, 1994), in the logit quantal response equilibrium (QRE; McKelvey and Palfrey, 1995), and in the experienced-weighted attraction model (EWA; Camerer and Ho, 1999).

The previous paragraph described the choice of a subject with level-1 strategic sophistication. Analyses of strategic sophistication in games, as summarized in Crawford et al. (2013), often find that many subjects can be described as level-1, level-2 (best response to level-1 behavior) and a type called D1 (Costa-Gomes et al., 2001), which eliminates dominated strategies for both players before responding to the uniform distribution of the opponent's play. We therefore also allow such heterogeneity of strategic sophistication.[32]

Players can also be heterogeneous in their responsiveness to the expected payoffs, represented by the parameter λ. For the range of payoffs in our game, estimates of λ often lie between 0.1 and 0.5.[33] Thus, in simulations we draw individual λ's from a uniform distribution on interval $[0.1, 0.5]$.[34]

To model the dynamics, we use the experience-weighted attraction (EWA) learning model (Camerer and Ho, 1999). In the EWA model, players form attractions for each strategy and then probabilistically select a strategy according to the logistic equation (1) based on these attractions in place of expected payoff. Attractions are updated based on experience, with those strategies that either performed or would have performed best receiving more reinforcement. The EWA model thus builds in the reinforcement of best response to historical observations, a property of the adjustment of strategies that was identified in the previous section as being partially present in the data.

More precisely, each strategy i in the EWA model has an attraction A_{it} at the end of period t. These attractions are updated according to

$$A_{it} = \frac{\varphi N_{t-1} A_{it-1} + (\delta + (1-\delta) I_{it}) \pi_{it}}{N_t}, \tag{2}$$

where π_{it} is the payoff that strategy i obtained (or would have obtained) against the opponent's strategy in round t, $I_{it} = 1$ if strategy i was used by the player at time t (and $I_{it} = 0$ if it was not used), and φ, δ are parameters. The experience count N_t is updated according to $N_t = \rho N_{t-1} + 1$, where ρ is a parameter. The initial attractions A_{i0} and the initial experience count N_0 are also parameters of the model. Again, to reflect heterogeneity, the

[32]In the double auction, bidding strictly above value or asking strictly below cost is dominated by setting bid equal to value or ask equal to cost.

[33]See e.g. Stahl and Wilson (1994), McKelvey and Palfrey (1995), Camerer and Ho (1999). More recently, Ivanov (2011) and Nunnari and Zapal (2016) estimated λ's in this range (including for individual subjects). In our setting, the range of payoffs is 100 (from -70 to 30 for the non-extreme values, and from -10 (or -5) to 90 (or 95) for the extreme vaues). Choice probabilities in equation (1) are invariant to shifts in payoffs but not to changes in their scale, thus λ's need to be adjusted to reflect the range of payoffs.

[34]For the mean value $\lambda = 0.3$, the average extent of strategic behavior of extreme-value traders in our game is 24-36 and of non-extreme-value traders is 6-13 (depending on whether all or only undominated strategies are used, on the setting and on the presence of naive traders). These values are higher than the ones actually observed in the experiment. We will therefore modify (some) players' utility function, as explained below.

parameters can be different for different players.

We do not estimate the parameters of the EWA model but use the values found reasonable in Camerer and Ho (1999) and calibrated to our setting. The previous section found that the best response to the full history of plays has a larger impact on choice than the immediate past. In the EWA model, the higher values of the experience discount parameters φ and ρ lead to longer memory being important. We thus take these parameters from the uniform distrubution on $[0.8, 1]$ (with the condition that $\rho \leq \phi$), which is consistent with estimates from Camerer and Ho (1999). The hypothetical payoff discount parameter δ reflects the trade-off between finding a best response and reinforcing the previous play. From the previous section, both appear to be important; we take δ uniformly distributed on $[0, 1]$. The initial experience count is set to $N_0 = 1/(1-\rho)$, so that the magnitude of experience count stays roughly constant. Finally, the initial attractions A_{i0} are set equal to reflect either level-1 behavior (equal to the expected payoff against the uniform distribution on the other trader's strategies) or level-2 (expected payoffs against the distribution that a level-1 player would play with the given λ), with equal probability. Within each level of strategic sophistication, with probability 0.65 a simulated subject eliminates dominated strategies for both players and adjusts expected payoffs accordingly.[35]

The model describes the attractions of strategies for a player who has to make one choice. In our setting, each player can be in four situations: low- or high-cost seller, and low- and high-value buyer. Each player thus has four sets of attractions. Only the attractions for the current role (buyer or seller) are updated in a given period based on equation (2). The model is also for a finite strategy space; we use the integer discretization from 0 to 100 as the strategy space for the simulations.[36]

To reflect some features of the observed data, we consider some modifications of the basic EWA model. Since 80% of choices in the experiment are in multiples of 5, we take the probability 0.8 that a choice determined by equation (1) is rounded to the nearest mutliple of 5. In the experiment, 76% of subjects played "naively" at least once, even though such behavior is weakly dominated. We thus take that with probability 0.75 a simulated subject has a linear (affine) cost of lying (in utility terms), with the fixed cost distributed uniformly on $[0, 4]$ (applies if ask is not equal to cost or bid not equal to value), and the marginal cost distributed uniformly on $[0, 0.2]$ for each unit of deviation from the cost (or the value).[37]

[35] At the first level of sophistication, this is similar to type D1 in Costa-Gomes et al. (2001). At the second level, this is similar to level-2 player who eliminates dominated strategies and best responds to D1. Although strategies with ask below cost or bid above value are played only 4% of the time in the experiment, 35% of subjects play them at least once. This is why we set this probability to 65%.

[36] In the experiment, choices could be entered as numbers with two digits after the decimal point. However, almost all choices were integers (97%); indeed, 80% of choices were in multiples of 5.

[37] These parameters are calibrated to get the proportion of naive play close to the one observed in the

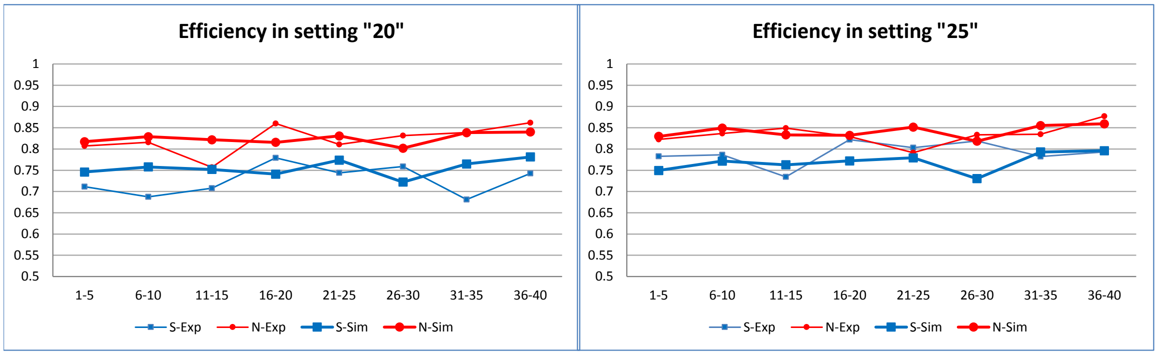

Figure 4: Efficiency in simulated sessions

In the post-experiment questionnaire, several subjects mentioned that they played "safe"; interpreting this as risk aversion, we use constant relative risk aversion (CRRA) utility function x^r to evaluate payoff, with r uniformly distributed on $[0.4, 1]$.[38]

Each simulation run models a matching group of 8 subjects, randomly matched for 40 rounds and each using the EWA learning as described above. For each simulated subject, parameters are drawn from the distributions specified above. The average results of 50 simulation runs for treatments S and N are presented in Figures 4 and 5, differentiated across settings "20" and "25".[39] Thick lines are the averages from the simulations while the experimental data are reproduced as thinner lines.

The average data from the simulations appear to fit the average data from the experiment quite well, both in terms of efficiency measures and in terms of the extent of strategic behavior. The realized efficiency in treatment S is noticeably lower than the expected efficiency in treatment N.[40] The extent of strategic behavior is only marginally higher for the extreme-value traders in the N treatment than in the S treatment, while for the non-extreme value traders there is little difference between the S and N treatments. It is also the case

data. The form of direct lying cost combines those discussed, e.g., in Ellingsen and Johannesson (2004), Kartik (2009) and Gneezy et al. (2018).

[38] More precisely, to avoid negative x and to preserve the range of utilities, the utility function is $\frac{100-0}{100^r-0^r}(x+70)^r$ for the non-extreme-value traders and $\frac{160-60}{100^r-60^r}(x+70)^r + \frac{60 \cdot 160^r - 60^r \cdot 160}{160^r-60^r}$ for the extreme-value traders. Ivanov (2011) uses a similar transformation. The average value of the risk-aversion parameter (0.7) corresponds to the one from (unincentivized) elicitation of risk attitudes in our communication sessions.

[39] In simulations of the N treatment, initial attractions based on the expected payoffs take into account that with probability 0.25 the opponent plays naively. Also, in each match simulated subjects play naive opponents with probability 0.25, as in the experiment. Attractions are still updated as described above.

[40] The human matches efficiency in treatment N (not shown to make the figure less cluttered) is only slightly lower than the realized efficiency in treatment S.

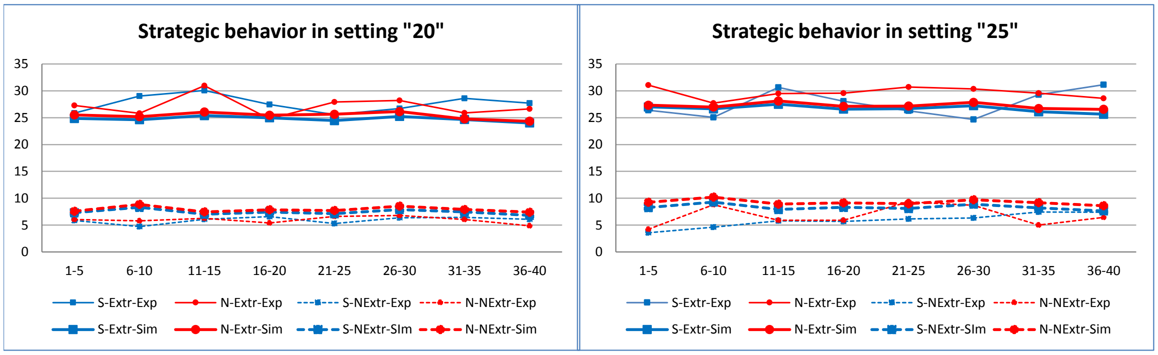

Figure 5: Strategic behavior in simulated sessions

that the proportions of naive play in the simulations are not too different from those in the experiment, and this proportion increases over time for non-extreme value traders.[41]

There are still some differences between the simulated and experimental data: the probabilities of trade in cases $\bar{v}-\bar{c}$ and $\underline{v}-\underline{c}$ are lower in simulations than in the experiment, and the probability of trade in case $\bar{v}-\underline{c}$ is higher. The simulations have a build-in symmetry across roles and do not distinguish between the BA and DM frames.[42] The fit is thus not perfect but the simulations demonstrate that a learning model can explain the behavior in the experiments to a large extent. In particular, simulations show that with noisy best response, even with learning, convergence to the efficient equilibrium is difficult, and that the presence of naive traders does not necessarily lead to more aggressive behavior and less efficiency, confirming what we observe in our experiments.

6 Conclusion

In this paper, we analyzed a bargaining situation with incomplete information in which a natural double-auction mechanism has an efficient equilibrium. However, with non-strategic naive traders, this efficient equilibrium disappears; thus, the presence of naive traders may decrease efficiency.

We ran an experiment to see whether the presence of naive traders would reduce ef-

[41]Depending on the setting ("20" or "25") and treatment ("S" or "N"), the proportions in simulations are 3-6% for extreme-value traders (compared with 5-9% in the experiment). For non-extreme-value traders, in simulations the proportions are 15-19% in the first five periods (20-27% in the experiment) but rise to 25-32% in the last five periods (compared with 26-39% in the experiment).

[42]A possible way to get the observed difference in the proportion of naive behavior between frames is to have a higher cost of lying for the DM frame or a larger probability that a player has a positive cost of lying.

ficiency but did not find this effect. The main reason for this is that the subjects did not play according to the efficient equilibrium in the treatment without naive traders. We then presented a learning model, with initial choices described by low levels of strategic sophistication and subsequent adjustment based on past experience, that replicates average measures of efficiency and behavior from the experimental data.

Taken at face value, our finding appears to contradict the previous literature (Radner and Schotter, 1989; Valley et al., 2002; McGinn et al., 2003; Ellingsen et al., 2009) which found that the constrained-efficient linear equilibrium of the double auction describes the data well even though the asymmetric information setting in those experiments (uniform distributions of value and cost) is more complex than our setting of two-point value/cost distributions. This apparent inconsistency can however be resolved by noting that the linear equilibrium of the double auction in the setting with identical uniform distributions of valuations can also arise from low levels of strategic sophistication (see footnote 9). Thus, play in the double auction may be better described by non-equilibrium models like level-k and learning dynamics in both settings.

The previous literature also showed that allowing pre-play communication generates efficiency gains that are greater than the theoretical maximum (e.g., Valley et al., 2002 and McGinn et al., 2003). However, we did not find a similar result in our setting: while allowing for pre-play communication did lead to an increase in efficiency, subjects did not come very close to the theoretical maximum of full efficiency. Our results thus suggest that caution is required in applying equilibrium concept to predict efficiency of private value bargaining outcomes.

The conclusions reached in this paper, in contrast to the previous literature, call for further experimental studies on private value bargaining. We need to understand better the nature of strategic behavior (whether it corresponds to equilibrium or to non-equilibrium models), as well as the possibility (or impossibility) of achieving full efficiency in practice, also (and perhaps especially) when it is theoretically feasible in equilibrium.

A Proofs

Proposition 1. *Suppose $\varepsilon = 0$. Strategies (Z_s^{str}, Z_b^{str}) are a fully efficient equilibrium if and only if $Z_s^{str}(\underline{c}) = Z_b^{str}(\underline{v}) = z_1$ and $Z_s^{str}(\bar{v}) = Z_b^{str}(\bar{c}) = z_2$, where*

$$\frac{2}{3}\underline{c} + \frac{1}{3}z_2 \leq z_1 \leq \underline{v} \quad \text{and} \quad \bar{c} \leq z_2 \leq \frac{2}{3}\bar{v} + \frac{1}{3}z_1. \tag{3}$$

Proof. In an equilibrium in which a (strategic) seller with value c trades with a positive probability, the ask cannot be strictly below the cost because the seller can get a higher payoff by setting the ask equal to the cost. Thus $c \leq Z_s^{str}(c)$. Similarly, $Z_b^{str}(v) \leq v$.

In a fully efficient equilibrium, $Z_s^{str}(\underline{c}) \leq Z_b^{str}(\underline{v})$ and $Z_s^{str}(\bar{c}) \leq Z_b^{str}(\bar{v})$. Combining with the inequalities from the previous paragraph,

$$\underline{c} \leq Z_s^{str}(\underline{c}) \leq Z_b^{str}(\underline{v}) \leq \underline{v} < \bar{c} \leq Z_s^{str}(\bar{c}) \leq Z_b^{str}(\bar{v}) \leq \bar{v}.$$

If $Z_s^{str}(\underline{c}) < Z_b^{str}(\underline{v})$, then $Z_c^{str}(\underline{c})$ can be increased to $Z_b^{str}(\underline{b})$ increasing seller's expected payoff. Thus $Z_s^{str}(\underline{c}) = Z_b^{str}(\underline{v}) = z_1 \leq \underline{v}$ and similarly $Z_s^{str}(\bar{c}) = Z_b^{str}(\bar{v}) = z_2 \geq \bar{c}$.

In equilibrium the seller with cost $\underline{c}$ cannot gain by increasing the ask to z_2. Thus $\frac{1}{2}\left(\frac{1}{2}z_1 + \frac{1}{2}z_2\right) + \frac{1}{2}z_1 - \underline{c} \geq \frac{1}{2}(z_2 - \underline{c})$, or $\frac{2}{3}\underline{c} + \frac{1}{3}z_2 \leq z_1$. A similar reasoning for the buyer with value $\bar{v}$ gives $z_2 \leq \frac{2}{3}\bar{v} + \frac{1}{3}z_1$. □

Proposition 2. *Suppose $\varepsilon \in (0,1)$. Strategies (Z_s^{str}, Z_b^{str}) are a fully efficient equilibrium if and only if $Z_s^{str}(\underline{c}) = Z_b^{str}(\underline{v}) = z_1$ and $Z_s^{str}(\bar{c}) = Z_b^{str}(\bar{v}) = z_2$, where*

$$\max\{A(\underline{c}, \underline{v}), B(z_2, \underline{c}, \underline{v}), C(z_2, \underline{c}, \underline{v}, \bar{v})\} \leq z_1 \leq D(\underline{c}, \underline{v}) \tag{4}$$

and

$$E(\bar{c}, \bar{v}) \leq z_2 \leq \min\{F(\bar{c}, \bar{v}), G(z_1, \bar{c}, \bar{v}), H(z_1, \bar{c}, \bar{v}, \underline{c})\}. \tag{5}$$

Proof. From the arguments in the proof of Proposition 1, in a fully efficient equilibrium, the inequalities $\underline{c} \leq z_1 \leq \underline{v} < \bar{c} \leq z_2 \leq \bar{v}$ hold. But now with a positive ε, we must have $\underline{c} < z_1 < \underline{v}$ and $\bar{c} < z_2 < \bar{v}$: if e.g. $\underline{c} = z_1$, then the low-cost seller does not gain anything when she trades at the price of z_1, and hence would be better off asking at least $\underline{v}$.

The seller with cost $\underline{c}$ gets in equilibrium the expected payoff

$$(1-\varepsilon)\left(\frac{1}{2}(z_1 - \underline{c}) + \frac{1}{2}\left(\frac{z_1 + z_2}{2} - \underline{c}\right)\right) + \varepsilon\left(\frac{1}{2}\left(\frac{z_1 + \underline{v}}{2} - \underline{c}\right) + \frac{1}{2}\left(\frac{z_1 + \bar{v}}{2} - \underline{c}\right)\right) \tag{6}$$

from setting ask equal z_1. From asking $\underline{v}$ the seller would get

$$(1-\varepsilon)\left(\frac{1}{2} \cdot 0 + \frac{1}{2}\left(\frac{\underline{v} + z_2}{2} - \underline{c}\right)\right) + \varepsilon\left(\frac{1}{2}\left(\frac{\underline{v} + \underline{v}}{2} - \underline{c}\right) + \frac{1}{2}\left(\frac{\underline{v} + \bar{v}}{2} - \underline{c}\right)\right). \tag{7}$$

The difference between payoffs (6) and (7) is non-negative if and only if $z_1 \geq A(\underline{c}, \underline{v})$. The seller would get the expected payoff

$$(1-\varepsilon)\left(\frac{1}{2} \cdot 0 + \frac{1}{2}(z_2 - \underline{c})\right) + \varepsilon\left(\frac{1}{2} \cdot 0 + \frac{1}{2}\left(\frac{z_2 + \bar{v}}{2} - \underline{c}\right)\right) \tag{8}$$

from setting ask equal z_2. The difference between payoffs (6) and (8) is non-negative if and only if $z_1 \geq B(z_2, \underline{c}, \underline{v})$. Finally, from asking $\bar{v}$ the seller would get

$$(1-\varepsilon)\left(\frac{1}{2}\cdot 0+\frac{1}{2}\cdot 0\right)+\varepsilon\left(\frac{1}{2}\cdot 0+\frac{1}{2}\left(\bar{v}-\underline{c}\right)\right). \tag{9}$$

The difference between payoffs (6) and (9) is non-negative if and only if $z_1 \geq C(z_2, \underline{c}, \underline{v}, \bar{v})$. All other deviations give a lower payoffs than one of the deviations considered above thus z_1 can be part of equilibrium if and only if $z_1 \geq \max\{A(\underline{c}, \underline{v}), B(z_2, \underline{c}, \underline{v}), C(z_2, \underline{c}, \underline{v}, \bar{v})\}$.

The buyer with value $\underline{v}$ gets in equilibrium the expected payoff

$$(1-\varepsilon)\left(\frac{1}{2}(\underline{v}-z_1)+\frac{1}{2}\cdot 0\right)+\varepsilon\left(\frac{1}{2}\left(\underline{v}-\frac{z_1+\underline{c}}{2}\right)+\frac{1}{2}\cdot 0\right) \tag{10}$$

from bidding z_1. Bidding $\bar{c}$ and z_2 cannot be profitable deviations. Bidding $\underline{c}$ would get the buyer the expected payoff

$$(1-\varepsilon)\left(\frac{1}{2}\cdot 0+\frac{1}{2}\cdot 0\right)+\varepsilon\left(\frac{1}{2}\left(\underline{v}-\underline{c}\right)+\frac{1}{2}\cdot 0\right). \tag{11}$$

The difference in payoffs (10) and (11) is non-negative if and only if $z_1 \leq D(\underline{c}, \underline{v})$.

The second line of inequalities in the proposition follows analogously from considerations of deviations from z_2. □

Proposition 3. *Suppose that* $\underline{c} = 10$, $\bar{c} = 70$, $\underline{v} = 30$, $\bar{v} = 90$.

i. If $\varepsilon = 0$, *there exists a unique fully efficient equilibrium* $Z_s^{str}(10) = Z_b^{str}(30) = 30$ *and* $Z_s^{str}(70) = Z_b^{str}(90) = 70$.

ii. If $\varepsilon \in (0, 1)$, *there is no fully efficient equilibrium.*

Proof. From Proposition 1, without naive traders, the conditions for the existence of the fully efficient equilibrium are $\frac{2}{3}\cdot 10 + \frac{1}{3}z_2 \leq z_1 \leq 30$ and $70 \leq z_2 \leq \frac{2}{3}\cdot 90 + \frac{1}{3}z_1$. The first condition can be satisfied only if $z_2 \leq 70$. Then from the second condition $z_2 = 70$ and, back from the first condition, $z_1 = 30$. Thus $Z_s^{str}(10) = Z_b^{str}(30) = 30$ and $Z_s^{str}(70) = Z_b^{str}(90) = 70$ is the unique efficient equilibrium in this setting.

Consider now $\varepsilon \in (0, 1)$. For the given parameter values, $B(z_2, \underline{c}, \underline{v}) = \frac{1}{3-\varepsilon}(z_2 + 20 - 30\varepsilon)$, $D(\underline{c}, \underline{v}) = \frac{1}{2-\varepsilon}(60 - 50\varepsilon)$ and $E(\bar{c}, \bar{v}) = \frac{1}{2-\varepsilon}(140 - 50\varepsilon)$. Inequality (4) in Proposition 2 implies that $B(z_2, 10, 30) \leq D(10, 30)$, or that $(2-\varepsilon)z_2 \leq 140 - 130\varepsilon + 20\varepsilon^2$. From inequality (5), $E(70, 90) \leq z_2$, or $140 - 50\varepsilon \leq (2-\varepsilon)z_2$. The two inequalities imply that $140 - 50\varepsilon \leq 140 - 130\varepsilon + 20\varepsilon^2$. This last inequality is equivalent to $20\varepsilon(4-\varepsilon) \leq 0$, which does not hold for any $\varepsilon \in (0, 1)$. Therefore there are no values of z_1 and z_2 that satisfy inequalities in Proposition 2. □

Proposition 4. *Suppose that* $\underline{c} = 5$, $\bar{c} = 70$, $\underline{v} = 30$, $\bar{v} = 95$.

i. If $\varepsilon = 0$, *there exist several fully efficient equilibria.*

ii. If $\varepsilon \in (0, 1)$, *a fully efficient equilibrium exists for* $\varepsilon \in (0, \bar{\varepsilon}]$, *where* $\bar{\varepsilon} = (11 - \sqrt{101})/5 \approx 0.190$, *and does not exist for* $\varepsilon \in (\bar{\varepsilon}, 1)$.

Proof. For $\varepsilon = 0$, the conditions for the fully efficient equilibrium from Proposition 1 are $\frac{2}{3} \cdot 5 + \frac{1}{3} z_2 \leq z_1 \leq 30$ and $70 \leq z_2 \leq \frac{2}{3} \cdot 95 + \frac{1}{3} z_1$. There are many values of z_1, z_2 that satisfy the two conditions, for example $z_1 = 30, z_2 = 70$ or $z_1 = 28$, $z_2 = 72$.

Consider now $\varepsilon \in (0, 1)$. Inequality (4) in Proposition 2 means that

$$z_1 \leq D(\underline{c}, \underline{v}) = \frac{1}{2 - \varepsilon}(60 - 55\varepsilon). \tag{12}$$

Since $G(z_1, \bar{c}, \bar{v}) = (z_1 + 190 - 70\varepsilon)/(3 - \varepsilon)$, inequality (5), combined with (12) implies $(3 - \varepsilon)(2 - \varepsilon)z_2 \leq 440 - 385\varepsilon + 70\varepsilon^2$. Since $E(\bar{c}, \bar{v}) = (140 - 45\varepsilon)/(2 - \varepsilon)$, inequality (5) also implies $140 - 45\varepsilon \leq (2 - \varepsilon)z_2$. Thus $(3 - \varepsilon)(140 - 45\varepsilon) \leq 440 - 385\varepsilon + 70\varepsilon^2$, which is equivalent to $0 \leq 4 - 22\varepsilon + 5\varepsilon^2$. This holds if $\varepsilon \leq \bar{\varepsilon} = (11 - \sqrt{101})/5 \approx 0.190$ and does not hold for $\varepsilon > \bar{\varepsilon}$. Therefore for $\varepsilon > \bar{\varepsilon}$ there is no z_1, z_2 that satisfy the inequailities of Proposition 2.

Take

$$z_1 = D(5, 30) = \frac{1}{2 - \varepsilon}(60 - 55\varepsilon) \quad \text{and} \quad z_2 = E(70, 95) = \frac{1}{2 - \varepsilon}(140 - 45\varepsilon). \tag{13}$$

By construction, $z_1 \leq D(5, 30)$. It holds that $A(5, 30) \leq z_1$ if $(2 - \varepsilon)(40 + 20\varepsilon) \leq (3 - \varepsilon)(60 - 55\varepsilon)$. This holds if $\varepsilon \leq \sqrt{129}/6 - 9/6 \approx 0.393$. For the given z_2, $B(z_2, 5, 30) = (160 - 115\varepsilon + 30\varepsilon^2)/((2 - \varepsilon)(3 - \varepsilon))$. Then $B(z_2, 5, 30) \leq z_1$ if $32 - 23\varepsilon + 6\varepsilon^2 \leq (3 - \varepsilon)(12 - 11\varepsilon)$, or $\varepsilon \leq (11 - \sqrt{101})/5 = \bar{\varepsilon}$. Finally, $C(z_2, 5, 30, 95) = (-100 + 275\varepsilon - 100\varepsilon^2)/((2 - \varepsilon)(3 - \varepsilon)) \leq z_1$ if $-20 + 55\varepsilon - 20\varepsilon^2 \leq (3 - \varepsilon)(12 - 11\varepsilon)$. This holds if $\varepsilon \leq 50/31 - 2\sqrt{191}/31 \approx 0.721$. Therefore for z_1 and z_2 in (13), inequalities in (4) of Proposition 2 are satisfied if $\varepsilon \leq \bar{\varepsilon}$. Analogous reasoning shows that inequalities in (5) are also satisfied. Thus z_1 and z_2 in (13) constitute a fully efficient equilibrium for any $\varepsilon \leq \bar{\varepsilon}$. □

B Sample Instructions (Treatment DM-N-20)

Experiment Instructions

Please read these instructions carefully. Please do not talk to other participants and remain quiet throughout the experiment. If you have a question, please raise your hand.

You are about to participate in a decision-making experiment. Your earnings at the end of the experiment will depend on the decisions you and other participants make. If you make careful decisions, you can earn money in addition to the £5.00 participation fee. At the end of the experiment you will be asked to fill a questionnaire and will be paid in private your earnings in the experiment, added to the £5.00 participation fee.

Description of the Experiment

The experiment will last 40 rounds. In each round you will make one decision as described below. The computer screen will give all the necessary information for your decision. After all participants have made their decisions in a round, you will get a payoff in Points and a new round will begin. Your earnings will be the sum of Points you get in all rounds, converted to pounds at the rate of £0.20 for every 10 Points.

Description of a round in the experiment

Matching and Roles

In each round all participants are randomly matched in pairs. Your match in any given round is independent of what happened in the previous rounds.

In each matched pair, one participant has the role of the **Seller** of a fictitious good and the other participant the role of the **Buyer** of the good. Some of you will be the Seller in the first 10 rounds, the Buyer in the next 10 rounds, the Seller again in the next 10 and the Buyer again in the last 10 rounds. Others will be the Buyer in the first 10 rounds, the Seller in the next 10, the Buyer in the next 10 and the Seller in the last 10 rounds. In each round the computer screen will show your role.

Costs and Values

The **Seller** has a **Cost** of producing the good. The Cost is determined each round randomly and can be either **10** or **70**, each with equal 50% chance. Before making the decision, the Seller will be told what his or her Cost is in a given round. The Buyer will not be told the Seller's Cost.

The **Buyer** has a **Value** for the good. The Value is determined each round randomly and can be either **30** or **90**, each with equal 50% chance. Before making the decision, the Buyer will be told what his or her Value is in a given round. The Seller will not be told the Buyer's Value.

Decisions

The **Seller** and the **Buyer** make one decision each, simultaneously and independently, not seeing the other's decision. The **Seller** makes a report about his or her Cost by choosing a **Reported Cost**. The **Buyer** makes a report about his or her Value by choosing a **Reported Value**. The Reported Cost and the Reported Value can be **any number** between 0 and 100 with at most two digits after the decimal point. They do not have to be equal to the actual Cost or Value.

After the decisions are made, it is determined for each pair whether the participants are matched with the computer instead of with each other. With a **25% chance** each of the participants in the pair is **matched with the computer** and with **75% chance** the participants stay **matched with each other**.

When **the Seller** is matched with the computer, the **computer's decision** is to choose **Reported Value = Buyer's Value**. When **the Buyer** is matched with the computer, the **computer's decision** is to choose **Reported Cost = Seller's Cost**.

Payoffs and Feedback

After determining whether you are matched with the computer or with the other participant, your payoff is calculated as follows. **If the Reported Cost is strictly above the Reported Value**, the good is **not traded**. **If the Reported Cost is equal or below the Reported Value**, the good is **traded** at the **Trading Price** in the middle between the Reported Cost and the Reported Value:

$$\text{Trading Price} = \frac{\text{Reported Cost} + \text{Reported Value}}{2}.$$

If the good is **not traded**, both the Seller and the Buyer get payoff **0**.

If the good is **traded** at some Trading Price, the payoffs are (in Points)

$$\begin{aligned}\text{Your payoff as Seller} &= \text{Trading Price} - \text{Seller's Actual Cost},\\ \text{Your payoff as Buyer} &= \text{Buyer's Actual Value} - \text{Trading Price}.\end{aligned}$$

At the end of each round, you will be told what the Reported Cost and the Reported Value in your match were (you will not be told whether you were matched with the computer or with the other participant). You will also be told whether there was trade and at what Trading Price, and what your payoff was in this round.

Examples

Suppose that the Buyer has Value 90 and chooses Reported Value 55. Suppose further that the matched Seller has Cost 10 and chooses Reported Cost 70.

With a 25% chance each of the participants in the pair is matched with the computer. For the Buyer, the decision of the computer is to choose Reported Cost = Cost = 10. Then there is trade at Trading Price $(10 + 55)/2 = 32.5$. The Buyer's payoff is $90 - 32.5 = 57.5$.

For the Seller, the decision of the computer is to choose Reported Value = Value = 90. The Trading Price for the Seller is $(70 + 90)/2 = 80$. The Seller's payoff is $80 - 10 = 70$.

With a 75% chance the participants stay matched between themselves. Since the Reported Cost 70 is above the Reported Value 55, there is no trade. The Buyer's and the Seller's payoffs are 0.

Suppose now that the Buyer's Value is 30 while the decision about the Reported Value is still 55. If Reported Cost = 10, either because of matching with the computer or because the Seller chose Reported Cost = 10, there will be trade at Trading Price 32.5. The Buyer's payoff will be $30 - 32.5 = -2.5$. The Buyer suffers a loss because the Reported Value 55 is above Buyer's Value 30. Similarly, the Seller can suffer a loss if the Reported Cost is below the Seller's Cost.

If you have any questions about the experiment, please raise your hand now.

Please pay attention now to the computer screens. You will be asked questions on your understanding of the instructions. You will also practise with decision and feedback screens. After everyone has answered the questions and finished the practice, the experiment will begin.

C Data and Tests

Non-parametric tests compare aggregate variables, such as efficiency, incidence of naive behavior, etc., across treatments on the matching group level. The tests below look at the data in Periods 21-40, where behavior is more informed than in earlier periods.

C.1 Efficiency and the Probability of Trade

C.1.1 Efficiency Comparison

Realized and expected efficiency

The table below presents the realized (treatment S) and expected (treatment N) efficiency measures, in each of the matching groups in all treatments, and the average efficiency measures for each treatment.

Periods 21-40	Setting "20"				Setting "25"			
	BA-S	BA-N	DM-S	DM-N	BA-S	BA-N	DM-S	DM-N
MG1	0.738	0.837	0.656	0.839	0.852	0.881	0.857	0.722
MG2	0.769	0.851	0.691	0.761	0.671	0.835	0.782	0.906
MG3	0.643	0.692	0.804	0.829	0.722	0.827		
MG4	0.852	0.894	0.632	0.935	0.853	0.856		
MG5	0.656	0.844						
MG6	0.887	0.880						
Average	0.750	0.829	0.704	0.845	0.784	0.851	0.828	0.803

Efficiency in each matching group in each treatment

The p-values of the two-sided Wilcoxon-Whitney-Mann rank-sum tests of the efficiency measure being equal across treatments are reported in the following table.

Efficiency level tests p-values Tests of $H_0 : \mu_{row} = \mu_{column}$						
	Setting "20"			Setting "25"		
	BA-N	DM-S	DM-N	BA-N	DM-S	DM-N
BA-S	0.262	0.336	0.286	0.248	0.355	0.643
BA-N	-	0.019**	0.670	-	0.643	1.000
DM-S	-	-	0.043**	-	-	1.000
** - significant at 5% level						

Efficiency comparison across treatments

There is no significant differences in the comparisons of pairs of treatments that differ only in frame (that is, comparisons between "BA-S" and "DM-S" and between "BA-N" and "DM-N"). Combining the observations across frames, the tests' results are

Efficiency level tests p-values Tests of $H_0 : \mu_{row} = \mu_{column}$		
	Setting "20" N	Setting "25" N
S	0.023**	0.337
** - significant at 5% level		

Efficiency comparison combining BA and DM frames

Human matching efficiency

In treatment N, we also look at the efficiency that would be achieved if matching was only between human subjects. The table below presents this efficiency measure (for treatment

S the numbers are the same as in the table in the previous subsection).

Periods 21-40	Setting "20"				Setting "25"			
	BA-S	BA-N	DM-S	DM-N	BA-S	BA-N	DM-S	DM-N
MG1	0.738	0.747	0.656	0.752	0.852	0.814	0.857	0.593
MG2	0.769	0.764	0.691	0.633	0.671	0.743	0.782	0.857
MG3	0.643	0.510	0.804	0.730	0.722	0.747		
MG4	0.852	0.846	0.632	0.898	0.853	0.782		
MG5	0.656	0.761						
MG6	0.887	0.816						
Average	0.750	0.735	0.704	0.758	0.784	0.773	0.828	0.710

Efficiency in human matches in each matching group in each treatment

The p-values of the Wilcoxon-Whitney-Mann rank-sum test of the efficiency across treatment being equal are reported in the following table.

Human matches efficiency level tests p-values Tests of $H_0 : \mu_{row} = \mu_{column}$						
	Setting "20"			Setting "25"		
	BA-N	DM-S	DM-N	BA-N	DM-S	DM-N
BA-S	0.873	0.336	0.831	1.000	0.355	1.000
BA-N	-	0.286	0.670	-	0.355	1.000
DM-S	-	-	0.387	-	-	1.000

Human matches efficiency comparison across treatments

There are no significant differences in any comparison. Aggregating across the BA and DM frames, the tests are

Human matching efficiency tests p-values Tests of $H_0 : \mu_{row} = \mu_{column}$		
	Setting "20" N	Setting "25" N
S	0.650	0.749

Human matching efficiency comparison combining BA and DM frames

No siginificant differences are detected either.

C.1.2 Regressions of the Probability of Trade

The tables below report the results of the regressions of the expected probability of trade ($ExpTrade$) and of the probability of trade if all subjects were matched between themselves ($HumTrade$), for different gains of trade. The expected trade variable $ExpTrade$ can take values strictly between 0 and 1 thus for it a simple regression is used. The variable $HumTrade$ is a binary variable; a probit regression is used and marginal effects reported. Explanatory variables are dummy variables denoting treatment ($NPresent$; 1 if treatment is N and 0 if S), frame (DM; 1 if frame is DM, 0 if BA), or parameter setting ($Val25$; 1 if setting "25", 0 if "20"), as well as $Period$ variable. The regressions are run for positive surplus situations and for Periods 21-40.

Determinants of expected trade probability in Periods 21-40 Dependent variable: $ExpTrade$			
Variable	$\bar{v} - \underline{c}$	$\bar{v} - \bar{c}$	$\underline{v} - \underline{c}$
$NPresent$	0.060** (0.027)	0.082** (0.039)	0.133** (0.050)
DM	0.001 (0.028)	−0.044 (0.045)	−0.002 (0.051)
$Val25$	0.053** (0.025)	−0.004 (0.041)	0.062 (0.050)
$Period$	0.001 (0.002)	0.002 (0.004)	0.003 (0.004)
$Constant$	0.817*** (0.067)	0.563*** (0.127)	0.395*** (0.127)
Observations	677	646	708
Clusters	32	32	32

Standard errors clustered by matching group in parentheses. *** - significant at 1% level; ** - significant at 5% level.

The expected probability of trade is higher in the N treatment than in the S treatment, while the other variables have (almost) no significant effect.

For the probit regression of the probability of trade in human matches the average marginal effects are

Determinants of trade probability in human matches in Periods 21-40 Dependent variable: $HumTrade$			
Variable	$\bar{v} - \underline{c}$	$\bar{v} - \bar{c}$	$\underline{v} - \underline{c}$
$NPresent$	0.010 (0.035)	−0.066 (0.043)	−0.035 (0.058)
DM	−0.002 (0.035)	−0.038 (0.050)	0.002 (0.061)
$Val25$	0.063* (0.036)	−0.010 (0.045)	0.073 (0.059)
$Period$	0.003 (0.002)	0.001 (0.005)	0.005 (0.004)
$(Constant)$	0.568 (0.395)	0.181 (0.370)	−0.435 (0.349)
Observations	677	646	708
Clusters	32	32	32

Standard errors clustered by matching group in parentheses. * - significant at 10% level.

There are (essentially) no significant effects of the variables on the probability of trade in human matches.

C.2 Individual Behavior

C.2.1 Incidence of "naive" behavior

Extreme values

The next table presents the incidence of "naive" behavior (the proportion of times the ask is set equal to the cost or the bid is set equal to the value) of the players with the extreme values ($\underline{c} = 10$ or $\underline{c} = 5$ for the seller and $\bar{v} = 90$ or $\bar{v} = 95$ for the buyer) in Periods 21-40.

Periods 21-40	Setting "20"				Setting "25"			
	BA-S	BA-N	DM-S	DM-N	BA-S	BA-N	DM-S	DM-N
MG1	0.010	0.176	0.000	0.000	0.136	0.095	0.019	0.115
MG2	0.060	0.081	0.176	0.067	0.117	0.024	0.083	0.150
MG3	0.011	0.040	0.029	0.156	0.013	0.099		
MG4	0.064	0.250	0.081	0.038	0.012	0.011		
MG5	0.062	0.063						
MG6	0.013	0.012						
Average	0.036	0.107	0.067	0.063	0.072	0.055	0.045	0.130

Naive behavior of players with extreme values in each treatment

The p-values of the Wilcoxon-Whitney-Mann rank-sum tests of these proportions being equal are reported in the following table.

Extreme values naive behavior proportions tests p-values Tests of $H_0 : \mu_{row} = \mu_{column}$							
	BA-N-20	DM-S-20	DM-N-20	BA-S-25	BA-N-25	DM-S-25	DM-N-25
BA-S-20	0.109	0.522	0.522	0.394	0.286	0.317	0.046**
BA-N-20	-	0.454	0.394	0.522	0.522	0.739	0.505
DM-S-20	-	-	0.885	1.000	1.000	1.000	0.355
DM-N-20	-	-	-	1.000	1.000	1.000	0.355
BA-S-25	-	-	-	-	0.564	1.000	0.355
BA-N-25	-	-	-	-	-	0.643	0.064*
DM-S-25	-	-	-	-	-	-	0.121
** - significant at 5% level; * - significant at 10% level							

Comparison of naive behavior with extreme values across treatments

There are no systematic differences across either setting ("20" or "25") or treatment ("S" or "N"). Aggregating across these dimensions, the comparison between frames ("BA" or "DM") gives

Naive behavior Tests of $H_0 : \mu_{row} = \mu_{column}$	
	DM
BA	0.599

Naive behavior with extreme values comparison between frames

Non-extreme values

The following table shows the incidence of "naive" behavior of the players with the non-extreme values ($\bar{c} = 70$ for the seller and $\underline{v} = 30$ for the buyer) in Periods 21-40.

Periods 21-40	Setting "20"				Setting "25"			
	BA-S	BA-N	DM-S	DM-N	BA-S	BA-N	DM-S	DM-N
MG1	0.485	0.327	0.247	0.630	0.371	0.263	0.396	0.375
MG2	0.065	0.384	0.256	0.306	0.422	0.182	0.443	0.400
MG3	0.083	0.198	0.546	0.446	0.050	0.169		
MG4	0.122	0.333	0.349	0.099	0.253	0.274		
MG5	0.380	0.210						
MG6	0.120	0.127						
Average	0.214	0.263	0.357	0.363	0.281	0.219	0.418	0.386

Naive behavior of players with non-extreme values in each treatment

The p-values of the Wilcoxon-Whitney-Mann rank-sum tests of these proportions being equal are reported in the following table.

Non-extreme values naive behavior proportions tests p-values Tests of $H_0 : \mu_{row} = \mu_{column}$							
	BA-N-20	DM-S-20	DM-N-20	BA-S-25	BA-N-25	DM-S-25	DM-N-25
BA-S-20	0.262	0.201	0.286	0.831	0.394	0.182	0.317
BA-N-20	-	0.286	0.522	0.670	0.394	0.046**	0.096*
DM-S-20	-	-	0.773	0.773	0.248	0.355	0.355
DM-N-20	-	-	-	0.387	0.248	1.000	1.000
BA-S-25	-	-	-	-	0.564	0.165	0.355
BA-N-25	-	-	-	-	-	0.064*	0.064*
DM-S-25	-	-	-	-	-	-	0.439
** - significant at 5% level; * - significant at 10% level							

Comparison of naive behavior with non-extreme values across treatments

The main difference is between BA and DM treatments, although not all disaggregated tests pick up this difference. The test on the aggregate level finds the following result:

Naive behavior	
	DM
BA	0.016**
** - significant at 5% level	

Comparison of naive behavior with non-extreme values between frames

C.2.2 Extent of strategic behavior

Extreme values

The next table looks at the extent of strategic behavior of the players with the extreme values ($\underline{c} = 10$ or $\underline{c} = 5$ for the seller and $\bar{v} = 90$ or $\bar{v} = 95$ for the buyer) in Periods 21-40. Only asks different from cost and bids different from value are considered.

Periods 21-40	Setting "20"				Setting "25"			
	BA-S	BA-N	DM-S	DM-N	BA-S	BA-N	DM-S	DM-N
MG1	26.93	28.13	27.66	30.71	25.21	21.86	27.53	45.48
MG2	32.27	21.81	31.87	33.47	36.09	30.18	23.48	21.25
MG3	31.76	37.37	25.39	26.75	33.41	25.50		
MG4	22.32	28.98	29.08	17.67	23.04	29.71		
MG5	28.13	23.04						
MG6	16.42	21.17						
Average	26.53	27.21	28.07	27.17	28.98	27.01	25.94	35.18

Strategic behavior of players with extreme values in each treatment

The p-values of the Wilcoxon-Whitney-Mann rank-sum tests comparing the extent of strategic behavior of players with extreme values across treatments are

Extreme values strategic behavior level tests p-values Tests of $H_0 : \mu_{row} = \mu_{column}$						
	Setting "20"			Setting "25"		
	BA-N	DM-S	DM-N	BA-N	DM-S	DM-N
BA-S	0.873	0.670	0.831	0.564	0.643	1.000
BA-N	-	0.394	0.831	-	0.643	1.000
DM-S	-	-	1.000	-	-	1.000
* - significant at 10% level						

Comparison of the extent of strategic behavior with extreme values across treatments

There are no significant differences in the comparisons. Neither there are siginificant differences on the aggregate level:

Extreme values strategic behavior Tests of $H_0 : \mu_{row} = \mu_{column}$					
Setting "20"			Setting "25"		
	N	DM		N	DM
S	0.762		S	0.873	
BA		0.537	BA		0.734

Aggregate comparison of strategic behavior of extreme value traders

Non-extreme values

The following table presents the extent of strategic behavior of the players with the non-extreme values ($\bar{c} = 70$ for the seller and $\underline{v} = 30$ for the buyer) in Periods 21-40, again only for asks different from cost and bids different from value.

Periods 21-40	Setting "20"				Setting "25"			
	BA-S	BA-N	DM-S	DM-N	BA-S	BA-N	DM-S	DM-N
MG1	6.707	3.977	-0.475	5.556	9.565	7.950	5.414	10.083
MG2	8.069	5.474	8.844	8.644	6.531	7.084	5.607	4.542
MG3	7.124	5.595	2.829	8.435	6.329	6.014		
MG4	9.875	10.729	7.348	5.041	7.039	8.000		
MG5	3.510	3.423						
MG6	3.808	6.058						
Average	6.706	5.664	4.769	6.907	7.353	7.157	5.502	7.620

Strategic behavior of players with non-extreme values in each treatment

The p-values of the Wilcoxon-Whitney-Mann rank-sum tests of the extent of strategic behavior of players with non-extreme values across treatments being equal are

Non-extreme values strategic behavior level tests p-values Tests of $H_0 : \mu_{row} = \mu_{column}$						
	Setting "20"			Setting "25"		
	BA-N	DM-S	DM-N	BA-N	DM-S	DM-N
BA-S	0.522	0.522	0.670	0.773	0.064*	1.000
BA-N	-	0.670	0.522	-	0.064*	1.000
DM-S	-	-	0.564	-	-	1.000
* - significant at 10% level						

Comparison of the extent of strategic behavior with non-extreme values across treatments

Although there are marginal significant differences of one treatment in setting "25", the tests on the aggregate level do not find the differences significant:

Non-extreme values strategic behavior Tests of $H_0 : \mu_{row} = \mu_{column}$					
Setting "20"			Setting "25"		
	N	DM		N	DM
S	0.880		S	0.423	
BA		1.000	BA		0.174

Aggregate comparison of strategic behavior of non-extreme value traders

Thus no significant systematic differences are detected across treatments and frames, either for extreme or for non-extreme value traders.

C.3 Sessions with Pre-Play Communication

There were 128 instances of play in the two sessions with pre-play communication. The following table reports the summary of the bid and ask behavior and outcomes, disaggregated by actual cost and value traders had:

Actual Value-Cost	Number of instances	Average Bid	Average Ask	One-price (Bid=Ask)	Trade (Efficiency)
90-10	49	42.35	36.89	35	46 (0.939)
90-70	33	72.12	79.39	26	26 (0.788)
30-10	29	20.86	27.83	14	18 (0.621)
30-70	17	17.06	83.35	0	0
Total	128			75	0.884

For 2 instances, the audio transcripts of conversations are lost. For the remaining 126 instances, the relevant content of conversations (exchange of value/cost information and agreement) is summarized in the table below (in one instance the recording is partial since it was cut before the conversation ended).

For each actual value/cost realization, the table reports how often each combination of value/cost was stated (the three last columns list instances where at least one trader did not explicitly state value or cost: e.g., "30-No" means that buyer stated 30 but seller did not state the cost; "None" means that neither trader stated value or cost). The table also reports, below each number, the number of times an agreement was reached (with the number of times the agreement was on one price in parenthesis) and the number of times trade happened.

Actual	No.	Stated Value-Cost						
Value-Cost		90-10	90-70	30-10	30-70	30-No	No-70	None
90-10	48	**10**	7	19	7	2	1	2
		9(8),9	7(7),7	19(19),19	5(1),5	2(2),2	1(1),1	1(1),2
90-70	32	0	**23**	0	7	0	0	2
			22(22),19		4(2),4			2(2),2
30-10	29	0	0	**19**	6	1	0	3
				17(14),14	4(1),1	0(0),0		0(0),3
30-70	17	0	1	0	**16**	0	0	0
			1(0),0		4(2),0			

The bold numbers in the table correspond to mutual truthful revelation, which is common in cases in which at least one trader has non-extreme value, but not when both traders have extreme values. Agreements are achieved in many cases but they do not always lead to trade: e.g. only 14 out of 17 agreements in the case of value 30 and cost 10 truthfully revealed led to trade.

References

Abeler, J., Nosenzo, D., Raymond, C. (2016) "Preferences for Truth-telling", CESifo working paper 6087.

Camerer, C., Ho, T.-H. (1999) "Experience-Weighted Attraction Learning in Normal Form Games", *Econometrica* 67, 827-874.

Chatterjee, K., Samuelson, L. (1987) "Bargaining with Two-sided Incomplete Information: An Infinite Horizon Model with Alternating Offers", *Review of Economic Studies* 54, 175-192.

Chatterjee, K., Samuelson, W. (1983) "Bargaining under Incomplete Information", *Operations Research*, 31, 835-851.

Costa-Gomes, M., Crawford, V., Broseta, B. (2001) "Cognition and Behavior in Normal-Form Games: An Experimental Study", *Econometrica* 69, 1193-1235.

Crawford, V.P., Costa-Gomes, M.A., Iriberri, N. (2013) "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications" *Journal of Economic Literature* 51, 5-62.

Crawford, V.P. (2016) "Efficient Mechanisms for Level-k Bilateral Trading", working paper, University of Oxford.

Ellingsen, T., Johannesson, M. (2004) "Promises, Threats and Fairness", *Economic Journal* 114, 397-420.

Ellingsen, T., Johannesson, M., Lilja, J., Zetterqvist, H. (2009) "Trust and Truth", *Economic Journal* 119, 252-276.

Feri, F., Gantner, A. (2011) "Bargaining or Searching for a Better Price? – An Experimental Study", *Games and Economic Behavior* 72, 376-399.

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

Fischbacher, U., Föllmi-Heusi, F. (2013) "Lies in Disguise: An Experimental Study on Cheating", *Journal of the European Economic Association* 11, 525-547.

Gneezy, U. (2005) "Deception: The Role of Consequences", *American Economic Review* 95, 384-394.

Gneezy, U., Kajackaite, A., Sobel, J. (2018) "Lying Aversion and the Size of the Lie", *American Economic Review* 108, 419-453.

Greiner, B. (2015) "Subject Pool Recruitment Procedures: Organizing Experiments with ORSEE", *Journal of the Economic Science Association* 1, 114-125.

Ivanov, A. (2011) "Attitudes to Ambiguity in One-Shot Normal Form Games: An Experimental Study", *Games and Economic Behavior* 71, 366-394.

Kartik, N. (2009) "Strategic Communication with Lying Costs", *Review of Economic Studies* 76, 1359-1395.

Kneeland, T. (2017) "Mechanism Design with Level-k Types: Theory and an Application to Bilateral Trade", WZB SP II 2017-303.

Leininger, W., Linhart, P., Radner, R. (1989) "Equilibria of the Sealed-bid Mechanism for Bargaining with Incomplete Information", *Journal of Economic Theory* 48, 63-106.

Lundquist, T., Ellingsen, T., Gribbe, E., Johannesson, M. (2009) "The Aversion to Lying", *Journal of Economic Behavior & Organization* 70, 81-92.

McGinn, K., Thompson, L., Bazerman, M. (2003) "Dyadic Processes of Disclosure and Reciprocity in Bargaining with Communication", *Journal of Behavioral Decision Making* 16, 17-34.

McKelvey, R.D., Palfrey, T.R. (1995) "Quantal Response Equilibria for Normal Form Games", *Games and Economic Behavior* 10, 6-38.

Myerson, R., Satterthwaite, M. (1983) "Efficient Mechanisms for Bilateral Trading", *Journal of Economic Theory* 29, 265-281.

Nunnari, S., Zapal, J. (2016) "Gambler's Fallacy and Imperfect Best Response in Legislative Bargaining", *Games and Economic Behavior* 99, 275-294.

Radner, R., Schotter, A. (1989) "The Sealed-bid Mechanism: An Experimental Study", *Journal of Economic Theory* 48, 179-220.

Saran, R. (2011) “Bilateral Trading with Naive Traders”, *Games and Economic Behavior* 72, 544-557.

Saran, R. (2012) “How Naiveté Improves Efficiency in Trading with Pre-Play Communication”, *Economic Letters* 117(1), 311-314.

Satterthwaite, M., Williams, S. (1989) “Bilateral Trade with the Sealed Bid k-Double Auction: Existence and Efficiency”, *Journal of Economic Theory* 48, 107-133.

Serra-Garcia, M., van Damme, E., Potters, J. (2011) “Hiding and Inconvenient Truth: Lies and Vagueness”, *Games and Economic Behavior* 73, 244-261.

Spiekermann, S., Grossklags, J., Berendt, B. (2001) “E-privacy in 2nd Generation E-commerce: Privacy Preferences versus Actual Behavior”, in: Proceedings of the 3rd ACM Conference on Electronic Commerce. ACM, New York, pp. 38-47.

Stahl, D.O., Wilson, P.W. (1994) “Experimental Evidence of Players’ Models of Other Players”, *Journal of Economic Behavior & Organization* 25, 309-327.

Valley, K., Thompson, L., Gibbons, R., Bazerman, M.H. (2002) “How Communication Improves Efficiency in Bargaining Games”, *Games and Economic Behavior* 38, 127-155.