



(In)efficiency in private value bargaining with naive players: Theory and experiment [☆]

Alex Possajennikov ^{a,*}, Rene Saran ^b

^a University of Nottingham, United Kingdom

^b University of Cincinnati, United States of America

ARTICLE INFO

JEL classification:

C72
C78
C91
D82

Keywords:

Bargaining with private values
Double auction
Efficiency

ABSTRACT

The paper investigates two-player double-auction bargaining with private values in a setting with discrete two-point overlapping distributions of traders' valuations. We characterize parameter settings in which there exists a fully efficient equilibrium, and 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. We conduct an experiment to test the theoretical possibility that the presence of naive traders can reduce efficiency. We find, however, that efficiency is not lower in the presence of naive traders. Subjects mostly set bid/ask prices strategically but they do not coordinate on a fully efficient equilibrium and the extent of strategic behavior is not different in the presence of naive players. We show that a learning model of noisy strategy adjustment explains the observed behavior better than other (equilibrium or non-equilibrium) models.

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 reveal their private information even in the presence of monetary incentives to misrepresent it.¹ Saran (2011) incorporates the possibility of naive traders in private value

[☆] We thank the British Academy (Grant SG101000) for financial support and Centre for Decision Research and Experimental Economics (CeDEx) at the University of Nottingham for access to laboratory facilities. The paper benefited from audience comments at various seminar and conference presentations. We also thank the editor and anonymous referees for their comments.

* Corresponding author at: School of Economics, University of Nottingham, University Park, Nottingham NG7 2RD, United Kingdom.

E-mail address: alex.possajennikov@nottingham.ac.uk (A. Possajennikov).

URL: <https://sites.google.com/view/alex-possajennikov> (A. Possajennikov).

¹ See, for example, Fischbacher and Föllmi-Heusi (2013), Gneezy et al. (2018), and Abeler et al. (2019) for evidence of this in simple one-person reporting decisions, and Gneezy (2005), Lundquist et al. (2009) and Serra-Garcia et al. (2013) 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).

<https://doi.org/10.1016/j.jebo.2023.10.003>

Received 25 November 2021; Received in revised form 29 September 2023; Accepted 5 October 2023

0167-2681/© 2023 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY-NC license (<http://creativecommons.org/licenses/by-nc/4.0/>).

bargaining to show that a mechanism designer who tailors the mechanism to the proportion of naive traders can improve efficiency. However, if the mechanism is fixed, then the impact of naive traders on efficiency in bilateral trading is ambiguous. The *direct effect* of introducing naive traders (i.e., the change in outcome while keeping the behavior of strategic traders constant) is to improve efficiency (because the naive traders always trade when it is efficient). But the *indirect effect* (i.e., the change in outcome when strategic traders react optimally to the presence of naive traders), could be such that it either enhances or reduces efficiency.²

In this paper, we further investigate possible effects of naive traders on efficiency and individual behavior in private value bargaining. 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. In this setting, we take as *naive* the behavior of setting the bid equal to the private value, or the ask equal to the private cost, directly revealing the private information.

We are specially interested in the theoretical possibility that the presence of naive traders can reduce efficiency in this environment, which will be the case if the indirect effect is negative and dominates the trivially positive direct effect. Intuitively, an increase in the proportion of naive traders in the double-auction mechanism dilutes the inherent tradeoff between probability and terms of trade. This is because a marginal increase in strategic behavior (i.e., shading bid below value or exaggerating ask above cost) against a *fixed belief* in the behavior of the strategic opponent generates a greater improvement in the terms of trade and a smaller negative impact on the probability of trade as the proportion of naive traders increases (for technical details, see footnote 8). Due to this change in incentives, the strategic traders could increase the amount of their shading/exaggeration when the proportion of naive traders increases, thereby having a negative indirect effect on efficiency.

The indirect effect is, however, more complex, taking into account the reactions and counter-reactions of strategic traders. To have a chance of observing negative indirect and, more importantly, total effects in equilibrium, we characterize parameter settings in which fully efficient equilibria exist without naive traders. However, as the proportion of naive traders increases, strategic traders have an incentive to increase their strategic behavior and the efficient equilibria disappear. Thus, if the traders coordinate on the most efficient equilibrium, then, in theory, the introduction of naive traders should reduce efficiency.

Our confidence in observing 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 suggested by Chatterjee and Samuelson (1983) is the most efficient (even though not all trades with positive surplus are achieved in it).³ Although there are many other equilibria of the double auction in that case (Leininger et al., 1989; Satterthwaite and Williams, 1989), experimental data typically conform to the linear equilibrium (Radner and Schotter, 1989; Valley et al., 2002; McGinn et al., 2003; Ellingsen et al., 2009).

Even in our simple discrete setting, the double-auction mechanism has multiple equilibria, both in the absence and presence of naive traders. In the absence of naive traders, the efficiency attained in equilibrium can range from zero to full efficiency. Similarly, the equilibrium efficiency has a wide range in the presence of naive traders, although it is strictly positive on the lower end and strictly less than full efficiency on the upper end. Depending on which equilibria are played, efficiency could be higher or lower as the proportion of naive traders increases. Hence, equilibrium theory per se does not provide a clear prediction regarding the effect of naive traders on efficiency in our setting. We conduct an experiment to bring more clarity on this issue.

In our experiment, we consider two treatments. In one (control) treatment, each subject always plays against another (randomly rematched) subject in the experiment, with one subject assigned the role of the buyer while the other the role of the seller. In the other treatment, 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.

In the standard double auction, traders set prices, but an equivalent game can also be framed as a direct mechanism.⁴ 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. Since such a framing can affect the share of traders playing naively, we also compare the performance of the double auction framed as a bid-ask setting and as a direct mechanism with the same allocation rule.

The experiment results are mixed, with no decrease in efficiency in the presence of naive traders. Efficiency in fact increases for some settings but this is solely due to the positive direct effect of artificial naive traders rather than due to a change in strategic behavior by human subjects. The subjects do not appear to coordinate on the fully efficient equilibrium in the setting without artificial naive players, and there is no clear effect of the change in the proportion of naive traders on behavior, either measured by the extent to which ask/bid differs from cost/value or by the proportion of naive plays by human participants: thus, the results

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.

³ 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.

⁴ 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.

point toward a zero indirect effect. The direct mechanism frame has an effect to make some traders behave naively more often (thus providing some evidence of aversion to direct lying), 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 the extent of strategic shading or exaggeration. Most subjects do behave strategically, but the extent of strategic behavior does not change significantly across different settings to influence the aggregate measures of efficiency. We find that subjects often follow what worked for them in the past, but they also respond to incentives to some extent as bids and asks adjust towards the best response to the observed history.

Since we minimize repeated-game effects in our experiment by randomly rematching subjects in each period, the fact that the observed history significantly affects individual behavior suggests that a learning model might be appropriate to explain behavior. We specifically consider the experience-weighted attraction (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).⁵ This model (which also takes into account possible disutility of deviation from private cost/value) indeed has a better fit to the data than static (equilibrium or non-equilibrium) models with a comparable number of parameters, or an alternative model with preferences for fairness (see Section 6.2). Simulations of the EWA model reproduce many 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.⁶ Our analysis suggests that non-equilibrium models like level- k and learning dynamics may be more useful predictors of play in the double auction than equilibrium and its efficiency properties.

The rest of the paper is organized as follows. In Section 2, we present the theoretical model and conditions that characterize the fully efficient equilibrium. In Section 3, we discuss the design of our experiment and provide our hypotheses. We present the results related to efficiency and average strategic behavior in Section 4. In Section 5, we show how history of play influences individual behavior. In Section 6, we present the dynamic model of learning and compare its fit with some other models. We conclude in Section 7 and collect proofs in the Appendix.

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} or 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 $\underline{c} < \underline{v} < \bar{c} < \bar{v}$.⁷

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 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 $Z_s^{ts} : \{\underline{c}, \bar{c}\} \rightarrow \mathbb{R}$ and $Z_b^{tb} : \{\underline{v}, \bar{v}\} \rightarrow \mathbb{R}$ 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 $q^{ts, tb}(\cdot, \cdot) : \{\underline{c}, \bar{c}\} \times \{\underline{v}, \bar{v}\} \rightarrow [0, 1]$ denote the probability of trade between the traders depending on their dispositions and values. Given that $\underline{c} < \underline{v} < \bar{c} < \bar{v}$, trade is efficient if it happens in all cases except when cost is \bar{c} and value is \underline{v} . Therefore trade is fully efficient if $q^{ts, tb}(\bar{c}, \underline{v}) = 0$, $q^{ts, tb}(\underline{c}, \underline{v}) = q^{ts, tb}(\underline{c}, \bar{v}) = q^{ts, tb}(\bar{c}, \bar{v}) = 1$ for all t_s and t_b .

For any $\varepsilon \in [0, 1)$, define

$$\begin{aligned} A(\underline{c}, \underline{v}) &= \frac{1}{3-\varepsilon} (2(1-\varepsilon)\underline{c} + (1+\varepsilon)\underline{v}); & B(z_2, \underline{c}, \underline{v}) &= \frac{1}{3-\varepsilon} (z_2 + 2\underline{c} - \varepsilon\underline{v}); \\ C(z_2, \underline{c}, \underline{v}, \bar{v}) &= \frac{1}{3-\varepsilon} (\varepsilon\bar{v} + 2(2-\varepsilon)\underline{c} - \varepsilon\underline{v} - (1-\varepsilon)z_2); \\ D(\underline{c}, \underline{v}) &= \frac{1}{2-\varepsilon} (2(1-\varepsilon)\underline{v} + \varepsilon\underline{c}); & E(\bar{c}, \bar{v}) &= \frac{1}{2-\varepsilon} (2(1-\varepsilon)\bar{c} + \varepsilon\bar{v}); \end{aligned}$$

⁵ More precisely, the initial attractions of the EWA model are set equal to the expected payoff against the uniform distribution of strategies of the other player (level-1).

⁶ That the model also fits initial choices is consistent with level- k model being a good predictor of initial responses in games (see Crawford et al., 2013 for a survey). Crawford (2021) and Kneeland (2022) apply the level- k theory to private value bargaining from the point of view of mechanism design.

⁷ 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.

$$F(\bar{c}, \bar{v}) = \frac{1}{3-\epsilon} (2(1-\epsilon)\bar{v} + (1+\epsilon)\bar{c}); \quad G(z_1, \bar{c}, \bar{v}) = \frac{1}{3-\epsilon} (z_1 + 2\bar{v} - \epsilon\bar{c});$$

$$H(z_1, \bar{c}, \bar{v}, \underline{c}) = \frac{1}{3-\epsilon} (\epsilon\underline{c} + 2(2-\epsilon)\bar{v} - \epsilon\bar{c} - (1-\epsilon)z_1).$$

The following proposition characterizes fully efficient equilibria of the double-auction game.

Proposition 1. Suppose $\epsilon \in [0, 1)$. 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 the Appendix. In a fully efficient equilibrium, the low-cost strategic seller and the low-value strategic buyer (and correspondingly the high-cost strategic seller and the high-value strategic buyer) agree on one price z_1 (and correspondingly z_2). Then there is also trade if the low-cost strategic seller meets the high-value strategic buyer, no trade when the low-value strategic buyer meets the high-cost strategic seller, and efficient trade if a strategic trader meets a naive opponent (between naive traders, trade is always efficient). The inequalities in the proposition ensure that the strategic traders do not find it profitable to deviate from z_1 and z_2 . Depending on the values $\underline{c}, \underline{v}, \bar{c}, \bar{v}$ and ϵ , the fully efficient equilibrium may or may not exist.

If $\epsilon = 0$, then the inequalities in Proposition 1 reduce to $\frac{2}{3}\underline{c} + \frac{1}{3}z_2 \leq z_1 \leq \underline{v}$ and $\bar{c} \leq z_2 \leq \frac{2}{3}\bar{v} + \frac{1}{3}z_1$. With $\epsilon > 0$, there are more deviations to consider than with $\epsilon = 0$, since a strategic trader may attempt to extract full surplus from her naive counterpart. Therefore, the conditions for the existence of fully efficient equilibria are more stringent for $\epsilon > 0$ than for $\epsilon = 0$. 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.

Strategic players generally find it profitable to act more strategically (i.e., increase their shading/exaggeration) against any given belief about the behavior of the strategic opponent as the proportion of naive players increases. This is because with an increase in ϵ , a marginal increase in shading/exaggeration has a smaller negative effect on the probability of trade and a greater positive effect on the terms of trade.⁸ The incentive to increase strategic behavior may lead to a negative indirect effect on efficiency that in equilibrium dominates the direct effect of an increase in ϵ , as the previous paragraph shows. Thus, we may observe a fall in efficiency as the proportion of naive players increases.

Irrespective of whether a fully efficient equilibrium exists or not, there are multiple equilibria in the double-auction game, both in pure and mixed strategies. At the lower end of efficiency, there is a “no-trade” equilibrium in which strategic traders do not trade between themselves. For example, both types of strategic buyer bidding \underline{c} and both types of strategic seller asking \bar{v} is an equilibrium for any ϵ if $2\bar{v} + \underline{c} \leq 3\bar{c}$ and $3\underline{v} \leq \bar{v} + 2\underline{c}$, as we have in our experiment settings. If $\epsilon = 0$, then such an equilibrium implies zero efficiency but if $\epsilon > 0$, efficiency is positive and increasing with ϵ since naive traders trade between themselves.

2.3. Settings in the experiment

We use the results from Proposition 1 to choose the parameter settings in the experiment. In the main treatment, we consider $\underline{c} = 10$, $\bar{c} = 70$, $\underline{v} = 30$, $\bar{v} = 90$ (coded as “20” by the difference $\underline{v} - \underline{c} = \bar{v} - \bar{c} = 20$). 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 2. Suppose that $\underline{c} = 10$, $\bar{c} = 70$, $\underline{v} = 30$, $\bar{v} = 90$.

- i. If $\epsilon = 0$, 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 $\epsilon \in (0, 1)$, there is no fully efficient equilibrium.

The proof is in the Appendix. It shows that for $\epsilon = 0$, there is only one set of values z_1 and z_2 that satisfy the conditions for the existence of a fully efficient equilibrium. 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 explains why full efficiency is 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. Similarly, the strategic high-cost seller would like to exaggerate her ask to get more from the naive high-value buyer. The

⁸ For instance, suppose the high-value strategic buyer believes that the strategic seller's strategy is given by some distribution of asks F_s , which has the density f_s . Then the high-value strategic buyer's expected payoff when she bids $z_b > \bar{c}$ is $(1-\epsilon) \int_0^{z_b} (\bar{v} - (z_b + z_s)/2) f_s(z_s) dz_s + \epsilon ((1/2)(\bar{v} - (z_b + \bar{c})/2) + (1/2)(\bar{v} - (z_b + \underline{c})/2))$. Hence, the effect of a marginal increase in shading (i.e., marginal reduction in z_b) equals $((1-\epsilon)F_s(z_b) + \epsilon)/2 - (1-\epsilon)(\bar{v} - z_b)f_s(z_b)$. The first term, $((1-\epsilon)F_s(z_b) + \epsilon)/2$, is the improvement in the terms of trade due to a reduction in the price of any trade. The second term, $-(1-\epsilon)(\bar{v} - z_b)f_s(z_b)$, is a negative effect (representing the loss at the margin) due to a decrease in the probability of trading. Thus, the improvement in the terms of trade is greater while the negative impact on the probability of trade is smaller as ϵ increases. This argument generalizes to other cases too.

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 efficient 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 experiment should reduce efficiency.

The knife-edge property of the fully efficient equilibrium when $\varepsilon = 0$ can make it difficult to coordinate on it because some types of traders earn zero surplus in this equilibrium. Also, in the experiment (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.⁹

We therefore consider another setting, $\underline{c} = 5$, $\bar{c} = 70$, $\underline{v} = 30$, $\bar{v} = 95$ (coded as “25” for $\underline{v} - \underline{c} = \bar{v} - \bar{c} = 25$). When $\varepsilon = 0$, this 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 3. 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 the Appendix. 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.¹⁰

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 (or from setting bid/ask different from their value/cost), a fully efficient equilibrium also remains.¹¹

Since there are multiple equilibria in the double auction (including those asymmetric between the seller and buyer roles), what equilibrium is played, if any, may depend on many factors. Also, players may not necessarily coordinate on an equilibrium, act strategically or have selfish preferences. The theory thus does not offer unambiguous prediction. Our choice of parameter settings, however, is such that it is possible that the indirect effect of an increase in the proportion of naive players leads to fewer trades in equilibrium, and this is what we try to see in the experiment.

3. Experiment design

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). We consider experimental treatments differing along the following dimensions: the absence or presence of naive traders, different parameter settings discussed above, and the framing of interactions. Our main treatment difference is in the artificially induced proportion of naive traders. In treatments labeled “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.¹²

The distributions of costs and values follow the two settings discussed in the previous section. One setting (“20”) involves $\underline{c} = 10$, $\underline{v} = 30$, $\bar{c} = 70$, $\bar{v} = 90$. 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 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 “25” with $\underline{c} = 5$, $\underline{v} = 30$,

⁹ Nevertheless, even in this case one can expect that the increase in ε leads to lower efficiency. In Section S1.1 of Supplementary Materials online, we show that in a class of simple symmetric mixed equilibria (which exist for reasonable values of ε , viz. $\varepsilon < 0.367$), maximum efficiency decreases monotonically with ε .

¹⁰ In Section S1.2 of Supplementary Materials online we show that maximum efficiency in this setting equals full efficiency for $\varepsilon \leq \bar{\varepsilon}$ and then decreases monotonically with ε in the class of simple symmetric mixed equilibria that exist for $\varepsilon < 0.411$.

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 faces a 0.75 probability of being matched with the human opponent and 0.25 probability of being matched with a naive computerized opponent. For each subject, the probability of meeting a naive opponent in the N treatment increases. This correlated matching avoids replacing subjects' bids with “naive” bids, but makes it impossible for two naive (computerized) traders to be matched, unlike in the theoretical model. To capture the effect of such matches and make efficiency measures comparable with the theoretical ones, we adjust measures of expected efficiency in the experiment as explained in Section 4.1.

$\bar{c} = 70$, $\bar{v} = 95$. 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 in the N treatment.¹³ Hence, if we assume coordination on a fully efficient equilibrium in treatment S, then we can expect efficiency to be lower in treatment N than S.¹⁴ We thus have our first main hypothesis:

Hypothesis 1. Efficiency in treatment S is higher than in treatment N.

Compared with a fully efficient equilibrium in treatment S, the presence of naive traders in treatment N in theory makes strategic sellers with high cost and strategic buyers with low value (traders with the *non-extreme* values) increase shading of bids or exaggeration of costs, since they can get a higher expected payoff by exploiting naive traders.¹⁵ Again, assuming coordination on a fully efficient equilibrium in treatment S, we have our second main hypothesis:

Hypothesis 2. There is a higher extent of strategic behavior (shading/exaggeration) by the traders with the non-extreme values in treatment N than in treatment S.

The effect of introduction of naive traders on the equilibrium behavior of strategic low-cost sellers and high-value buyers (traders with the *extreme* values) is less clear cut. For example, in the unique fully efficient equilibrium in setting “20” with $\epsilon = 0$, 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. With $\epsilon = 0.25$, 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, the low-cost seller now acts less strategically whereas all other types of traders act more strategically.¹⁶ Due to this ambiguous effect, we do not state a formal hypothesis comparing the extent of strategic behavior by the traders with extreme values.

Our main hypotheses presume coordination on a fully efficient equilibrium in treatment S. In general though, there are multiple equilibria in the double auction and, without further assumptions/refinements, equilibrium theory does not offer an unambiguous prediction, as discussed in the last section. Beyond equilibrium, at a more basic level, an increase in the probability of naive traders has a positive direct effect on expected efficiency for the trivial reason that now a greater proportion of traders – the naive ones – are not shading/exaggerating their bids/asks. If the strategic traders react “instinctively” to the increase in naive behavior (i.e., while assuming that the behavior of the strategic opponent is held constant), then they have a greater incentive to shade/exaggerate because, as mentioned in the last section, a greater proportion of naive traders dilutes the tradeoff between probability and terms of trade. The resulting indirect effect could increase strategic behavior (from both non-extreme and extreme value traders), which could cause efficiency to decrease in matches where both traders are strategic. Thus, to the extent that players are prone to reacting instinctively, we can expect (at least some) human traders to behave more strategically with a higher proportion of naive traders, possibly leading to lower efficiency in human matches in treatment N, which is consistent with our hypotheses. Statistically, the hypotheses can be tested and average treatment effects identified by comparing efficiency and strategic behavior in the two treatments (provided that the assignment of subjects to treatments is random, as is the case in our experiment).

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.¹⁷ Framing the interaction as a direct mechanism makes it clearer what revealing one's cost or value is, and therefore may affect how subjects feel about exaggerating their cost or shading their value. 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 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.

¹³ For reasonable beliefs about ϵ being close to 0.1 (based, for example, on Valley et al., 2002, who find that about 10% of subjects make ask/bids equal to their cost/value in their no-communication treatment), the N treatment implies the increase in the probability of meeting a naive trader by $0.75 \cdot 0.1 + 0.25 \cdot 0.1 = 0.225$.

¹⁴ Even if beliefs about naive traders are not necessarily $\epsilon = 0$ in treatment S, efficiency still can be expected to decrease monotonically with ϵ , as we demonstrate in Section S1 of Supplementary Materials online, using simple symmetric mixed equilibria.

¹⁵ Even if beliefs about naive traders are not necessarily $\epsilon = 0$ in treatment S, we can expect strategic behavior (shading/exaggeration) of traders with the non-extreme values to increase monotonically with ϵ , as we demonstrate in Section S1 of Supplementary Materials online, using simple symmetric mixed equilibria.

¹⁶ 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, the high-value buyer acts less strategically whereas all other types of traders act more strategically.

¹⁷ The full set of instructions is available in the Experiment Instructions file online.

Table 1
Description of the experimental sessions.

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

Note: BA - bid/ask frame; DM - direct mechanism frame; S - treatment with $\epsilon = 0$, N - treatment with $\epsilon = 0.25$; 20 - setting “20”, 25 - setting “25”.

In total, there are 16 sessions (32 matching groups) with 272 subjects, who all passed the control questions on understanding the instructions.

A session lasts 40 periods. In each period, a subject in one role is randomly matched 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.¹⁸ After the matching, each subject's value or cost is drawn randomly and independently. The subjects then make their decisions by inputting a number between 0 and 100 with at most two digits after decimal point¹⁹ and their payoffs are calculated.²⁰ 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).²¹ 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 in March 2012. 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.²² Sessions were assigned to treatments randomly, making the assignment of subjects to treatment random and allowing identification of average treatment effects. Together with reading the instructions, answering control questions, and filling in the post-experiment questionnaire with demographic information, 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.

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 N, this measure of efficiency is affected by the correlated matching with computerized opponents, which makes both subjects in a match play against the computerized traders. However, in the theoretical model, the dispositions (being naive or strategic) of the players are independent rather than correlated.

To have a measure of efficiency that is comparable with the theoretical model, we calculate *expected probability of trade* and *expected efficiency* as if the dispositions were independently determined. In each matched pair of human subjects, given the realized cost c and value v and the subjects' bid b and ask a , the expected probability of trade (for given ϵ) is $PT_e = (1 - \epsilon)^2 \cdot I_{b \geq a} + (1 - \epsilon)\epsilon \cdot I_{b \geq c} + \epsilon(1 - \epsilon) \cdot I_{v \geq a} + \epsilon^2 \cdot I_{v \geq c}$, where I_X is the indicator variable for condition X ($I_X = 1$ if condition X is satisfied and $I_X = 0$ if condition X is not satisfied). The expected surplus for this match is $PT_e \cdot (v - c)$. These calculations allow to capture the direct positive effect of naive traders on efficiency, which is not available with the correlated matching.²³ The available surplus in a pair is $(v - c) \cdot I_{v \geq c}$. The overall expected efficiency is the ratio of the total expected surplus to the total available surplus over all matched pairs.

Another measure of efficiency we use is 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 probability of trade* in a match with given bid b and ask a is $PT_h = I_{b \geq a}$. The surplus that would have been realized is $PT_h \cdot (v - c)$. The *human efficiency* is the ratio of this surplus summed up over all matched pairs to the total available surplus over the pairs.

In treatment S ($\epsilon = 0$), the expected efficiency measure and the human efficiency measure coincide (and they are also equal to the realized efficiency). The theoretical analysis of the situations played out in the experiment shows that if the strategic traders

¹⁸ 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.


¹⁹ 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.

²⁰ In the N treatment, it is first determined whether the subjects in a pair are rematched against computerized opponents and then payoffs are calculated.

²¹ Recall that our games have multiple equilibria. Having multiple rounds with feedback provides an opportunity to coordinate on a particular equilibrium, while random matching makes the game look closer to a one-shot game.

²² Overall, 54% of the experiment participants were female; the average age was 20.0 years.

²³ For example, if the buyer's value is 30 and the seller's cost is 10, and the subject in the buyer's role bids 20 while the subject in the seller's role asks 75, then the expected probability of trade with $\epsilon = 0.25$ and independent dispositions is $(0.75)^2(0) + (0.75)(0.25)(1) + (0.25)(0.75)(0) + (0.25)^2(1) = 0.25$. The expected surplus is $0.25 \cdot 20 = 5$. In the experiment, because of the correlated matching, with probability 0.75 realized surplus is 0, and with probability 0.25, realized surplus is 20 for the buyer and 0 for the seller. Averaging over the pair, the realized surplus is $0.25(20 + 0)/2 = 2.5$, smaller than the expected surplus.



Note: Blocks of periods on the horizontal axis; fraction of surplus realized on the vertical axis, aggregated across BA and DM frames. N-e: expected efficiency in treatment N; N-h: human efficiency in treatment N, as defined in text (in treatment S, expected and human efficiency are the same).

Fig. 1. Efficiency measures in treatments S and N. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

Table 2
Probabilities of trade for different trade surpluses.

Trade surplus	Setting "20"		Setting "25"	
	S	N	S	N
	$PT_e = PT_h$	$PT_e (PT_h)$	$PT_e = PT_h$	$PT_e (PT_h)$
$\bar{v} - \underline{c}$	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)
$\underline{v} - \bar{c} (< 0)$	0.015	0.018 (0.023)	0.008	0.043 (0.050)

Note: Averages from all rounds and across BA and DM frames. Expected and human probabilities of trade PT_e and PT_h are as defined in text.

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) than with a sizable proportion of naive traders (treatment N). 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.

Fig. 1 presents the comparison of efficiency in treatments S and N. It shows the time series in the experiment of the three efficiency measures – realized efficiency in treatment S labeled “S” (blue, lighter line), expected efficiency in treatment N labeled “N-e” (red, darker line), and human efficiency in treatment N labeled “N-h” (dotted line) – averaged over blocks of five periods. The efficiency measures are separated across the “20” and “25” settings since we expect efficiency to be different in setting “25”, where positive trade surpluses are larger. They are combined for the BA and DM frames, since the efficiency measures are very similar between the frames,²⁴ and the regression results below do not find significant differences between the frames.

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 efficiency in treatment S, especially in setting “20”. There appears to be little difference between the treatments when comparing the efficiency in treatment S with the human efficiency in treatment N.

The efficiency measures presented in Fig. 1 are aggregate measures, combining 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 negative trade surplus ($\underline{v} - \bar{c}$). Table 2 shows the expected and human average probabilities of trades for the different surplus situations (combined across BA and DM frames).²⁵

From Table 2, the expected probability of trade is higher in treatment N than the probability of trade in treatment S for all trade surpluses (even for the negative one $\underline{v} - \bar{c}$, although such trades happen only in less than 3% of all cases of negative surplus).²⁶ However, the human probability of trade in treatment N is very similar to the probability of trade in treatment S for all gains of trade.

To test formally the effect of the treatments on the probability of trade, and thus on efficiency, we regress the probabilities of trade on the treatment variables in Table 3. In columns I-II of the table, the dependent variable is the expected probability of trade (PT_e); in columns III-IV, it is the probability of trade if all subjects were matched between themselves (PT_h). The variable PT_e can take values 0, 1, and also in between (see footnote 23), while PT_h is a binary variable. Since we are interested in treatment average effects rather than predicted probability, we use ordinary least square regressions for both.²⁷ Explanatory variables are dummy

²⁴ The measures for each matching group in each treatment are presented in Section S2.1 of Supplementary Materials online.

²⁵ In the S treatment ($\epsilon = 0$), the probabilities coincide (and equal the realized proportion of trades); thus only one number is presented for treatment S.

²⁶ 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.

²⁷ Average treatment effects (i.e., changes in the predicted value of the dependent variable as a treatment variable changes from 0 to 1, averaged over all sample values of the other independent variables) from the probit regression of PT_h , reported in Section S2.2 of Supplementary Materials online, are similar to those found in the linear regression.

Table 3
Effects of treatment variations on the probability of trade.

Variables	Dependent variables			
	Expected trade probability PT_e		Human trade probability PT_h	
	I	II	III	IV
Treatment N	0.088*** (0.020)	0.081*** (0.021)	−0.004 (0.024)	−0.012 (0.025)
Frame DM	−0.012 (0.022)	−0.008 (0.023)	−0.009 (0.027)	−0.005 (0.029)
Setting “25”	0.041* (0.022)	0.031 (0.023)	0.044 (0.026)	0.035 (0.028)
Period		0.001** (< 0.001)		0.001*** (< 0.001)
$\bar{v} - \underline{c}$		0.371*** (0.020)		0.430*** (0.018)
$\bar{v} - \bar{c}$		0.084*** (0.024)		0.094*** (0.026)
$\underline{v} - \bar{c}$		−0.517*** (0.025)		−0.430*** (0.023)
Constant	0.464*** (0.023)	0.465*** (0.030)	0.462*** (0.026)	0.420*** (0.030)
Observations		5440		5440
R^2	0.01	0.48	< 0.01	0.38

Note: Treatment S, frame BA and setting “20” are the base categories in all columns; trade surplus $\underline{v} - \underline{c}$ is the base category in columns II and IV. Standard errors clustered by 32 matching groups in parentheses. *** - significant at 1% level; ** - significant at 5% level; * - significant at 10% level.

variables denoting treatment (Treatment N: 1 if treatment is N and 0 if S), frame (Frame DM: 1 if frame is DM, 0 if BA), or parameter setting (Setting “25”: 1 if setting is “25”, 0 if “20”). In a fuller model, we also include the Period variable and dummy variables for available surplus ($\bar{v} - \underline{c}$, $\bar{v} - \bar{c}$, $\underline{v} - \bar{c}$; surplus $\underline{v} - \underline{c}$ acts as the base category).²⁸

The regressions show that the probability of trade, and thus efficiency, is not affected by frame (BA or DM), and hardly by setting (“20” or “25”). The probabilities of trade increase slightly over time but this effect is very small. Most significantly, the expected probability of trade, and thus the expected efficiency, is higher in treatment N than in treatment S, but the probability of trade in matches between subjects is the same between the treatments.²⁹ The introduction of naive traders thus raises the expected probability of trade but does not change the probability of trade in matches between human subjects.

Result 1. (Efficiency)

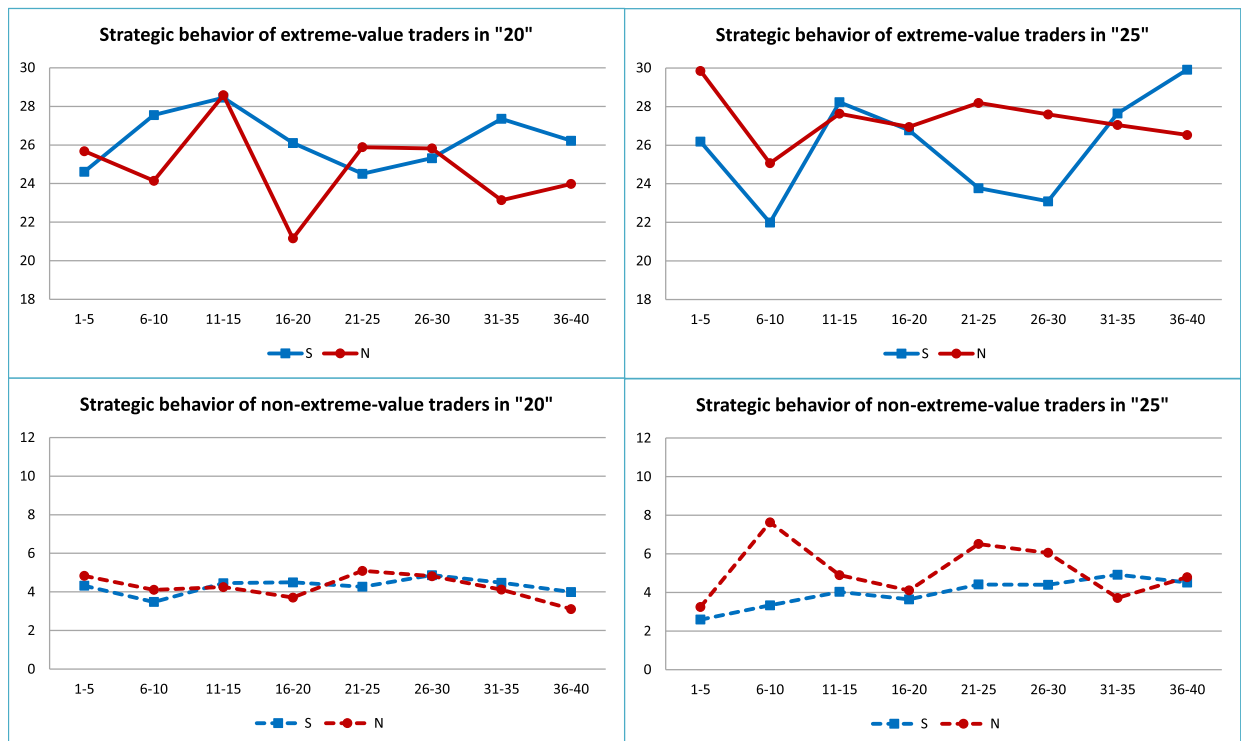
- Expected efficiency is higher in treatment N than in treatment S.
- Human efficiency is not significantly different between treatments N and S.
- There are no significant differences in efficiency between BA and DM frames.

Thus our hypotheses related to efficiency are not confirmed. Recall that if subjects were to coordinate on the fully efficient equilibrium in treatment S, then efficiency in treatment S will be higher than both the expected and human efficiencies in treatment N. The results, however, go in the opposite direction: the expected efficiency in treatment N is higher than the efficiency in treatment S in our experiment. Since we find that the human efficiency in treatment N is not different from the efficiency in treatment S, the greater value of the expected efficiency in treatment N compared to the efficiency in treatment S is mostly due to the trivial direct effect, i.e., the additional trades (that would have happened if traders’ dispositions were determined independently) with and between computerized naive traders rather than between human subjects.

An obvious reason for the failure to get higher efficiency in the S treatment compared to the N treatment is that subjects are not playing a fully efficient equilibrium in treatment S. The absence of a significant difference between efficiency in the S treatment and efficiency in matches of human subjects in the N treatment points also to the possibility of a zero indirect effect, that is, subjects do not behave more strategically in treatment N compared with treatment S. We analyze strategic behavior in more detail in the next subsection.

²⁸ The regressions thus include observations with negative trade surplus $\underline{v} - \bar{c}$. If such observations are excluded, the significance of the coefficients remains the same and their magnitudes are similar, as reported in Section S2.2 of Supplementary Materials online. Regressions on subsamples of situations with different positive surpluses, also reported in Section S2.2 of Supplementary Materials online, show that the effects of the treatment variables are similar in situations of different surpluses.

²⁹ Non-parametric tests at the matching group level, reported in Section S2.1 of Supplementary Materials online also confirm this, especially for setting “20”. So do regressions using all possible matches in a period in a matching group (rather than only realized matches as in Table 3), as reported in Section S2.2 of Supplementary Materials online.



Note: Blocks of periods on the horizontal axis; average strategic behavior ($v - b$ and $a - c$) on the vertical axis, aggregated across BA and DM frames. Top panels: extreme values (\underline{c} and \bar{v}); bottom panels: non-extreme values (\underline{v} and \bar{c}).

Fig. 2. Extent of strategic behavior in treatments S and N.

4.2. Strategic behavior

The previous subsection looked at the efficiency measures in the experimental sessions, which are derived from players' actual bids (b) and asks (a) (or statements about their values (v) or costs (c) in frame DM). In this subsection, we analyze players' behavior in more detail. In particular, we are interested by how much their bids and asks differ from their value or cost (the extent of strategic shading/exaggeration). We term this difference ($v - b$ for buyers and $a - c$ for sellers) *strategic behavior*.

Fig. 2 shows the time series of the average amount of shading and exaggeration in the experiment. The data are presented separately for extreme- (top panels) and non-extreme-value traders (bottom panels), and they are also separated for the "20" (left panels) and "25" (right panels) settings, since in setting "25" the scope for ask/bids further away from cost/value is larger. Behavior is further disaggregated between S (labeled 'S-') and N (labeled 'N-') treatments (while aggregated across BA and DM frames).³⁰

Unsurprisingly, traders with the extreme values shade or exaggerate much more (on average, 25.97) than traders with the non-extreme values (4.35 on average). From the figure, there appears to be little difference in the average extent of strategic behavior between treatments S and N. 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 is not fulfilled. The figure also confirms that the play is not that close to the fully efficient equilibrium: in "20" setting, the non-extreme-value traders on average shade value/exaggerate cost by 5 whereas the extreme-value traders shade value/exaggerate cost by more than 20.

To check these observations from the figure more formally, we use regression analysis. In the regressions, the dependent variable is strategic behavior as defined above ($a - c$ or $v - b$). Among the explanatory variables, Treatment N, Frame DM, and Setting "25", as before, are the dummy variables describing the treatment, frame, and setting. We run one set of regressions, separately for the subsamples of observations with the extreme and with the non-extreme values, with only these variables. In another set, we also include the dummy variable Seller (equal to 1 if a subject is a seller in period t and 0 if the subject is a buyer) and the variable Period. The results of the regressions are presented in Table 4.

³⁰ As regressions below show, there is no significant differences in strategic behavior between frames. The averages of strategic behavior in all matching groups are presented in Section S3.1 of Supplementary Materials online.

Table 4
Effects of treatment variations on strategic behavior.

Variables	Dependent variable: Strategic behavior			
	Extreme		Non-extreme	
	I	II	III	IV
Treatment N	−0.391 (1.812)	−0.373 (1.815)	0.383 (0.713)	0.382 (0.714)
Frame DM	1.192 (2.069)	1.115 (2.081)	−1.361* (0.801)	−1.363* (0.801)
Setting “25”	1.164 (2.086)	1.119 (2.094)	0.192 (0.768)	0.191 (0.768)
Seller		5.668*** (1.182)		−0.064 (0.435)
Period		−0.006 (0.037)		0.006 (0.015)
Constant	25.28*** (1.758)	22.55*** (1.758)	4.606*** (0.700)	4.523*** (0.700)
Observations	5411		5469	
R ²	< 0.01	0.02	< 0.01	< 0.01

Note: Strategic behavior is $a - c$ or $v - b$. Extreme: values \underline{c} and \bar{v} ; non-extreme: values \underline{v} and \bar{c} . Treatment S, frame BA, and setting “20” are the base categories in all columns; Buyer is the base category in columns II and IV. Standard errors clustered by 32 matching groups in parentheses. *** - significant at 1% level; * - significant at 10% level.

The regressions confirm that the average strategic behavior is similar across treatments, frames, and settings: none of the dummy variables for those are significant.³¹ Sellers appear to be more strategic than buyers when having extreme values, while somewhat less strategic with non-extreme values.³² No significant time trend is detected in the average strategic behavior.

Recall that we consider behavior naive if bid is set equal to value or ask equal to cost (or, in frame DM, actual value or cost is reported). This corresponds to strategic behavior being equal to 0 and is captured by variable Naive, equal 1 if subject i played naively in period t and 0 otherwise. Even though the average strategic behavior in the experiment is clearly different from 0, subjects sometimes play naively: 7% of choices of traders with extreme values and 28% of choices with non-extreme values are “naive”.³³ Table 5 reports the results of linear regressions (again run separately for subsamples of extreme and non-extreme value traders) for dependent variable Naive.³⁴

The regressions find that there is significantly more naive play in frame DM (36%) than in frame BA (24%) when traders have non-extreme values (\underline{v} or \bar{c}).³⁵ With extreme values, sellers play naively less often than buyers, while with non-extreme values the opposite is true. There is a significant increase over time in naive play by traders with non-extreme values (from 25% in the first 10 periods to 32% in the last 10).

Summarizing the findings about the average strategic behavior so far:

Result 2. (Average strategic behavior)

- There are no significant differences in the extent of strategic behavior across treatments and frames; traders with non-extreme values play naively quite often, and more often in frame DM than in frame BA.
- With extreme values, sellers play more strategically than buyers but with non-extreme values, there is no difference in strategic behavior between sellers and buyers.

³¹ Non-parametric tests at the matching group level, reported in Section S3.1 of Supplementary Materials online, find the same result. Regressions combining observations across subsamples, reported in Section S3.3 of Supplementary Materials online, also produce the same results. Standard errors in the reported regressions are clustered by matching group, since a subject's behavior can depend on what happens in the whole matching group, not only on the subject's own characteristics and behavior. However, as reported in Section S3.3 of Supplementary Materials online, the significance of the variables is the same with standard errors clustered at the individual level.

³² This difference is one explanation for the different probabilities of trade noted earlier: high-value buyers and high-cost sellers trade more often because they are less strategic than low-cost sellers and low-value buyers. While this difference between roles is somewhat surprising (the game is symmetric and all subjects played an equal number of times in each role), Radner and Schotter (1989) also find a difference in buyer and seller behavior in their experiment on symmetric double auction.

³³ The large difference between these proportions shows that few subjects consistently played naively for both extreme and non-extreme values. Since playing naively with an extreme value has very little strategic sense, 7% can be taken as an indicator of naturally occurring naive players. This proportion is close to the 10% proportion of “naive” (ask/bid equal cost/value) plays in Valley et al. (2002) and can further justify that reasonable beliefs of subjects about ε (not far from 10%) allow for *a priori* hypothesis about efficiency comparison between the S and N treatments. The proportions of naive play in each matching groups are reported in Section S3.2 of Supplementary Materials online.

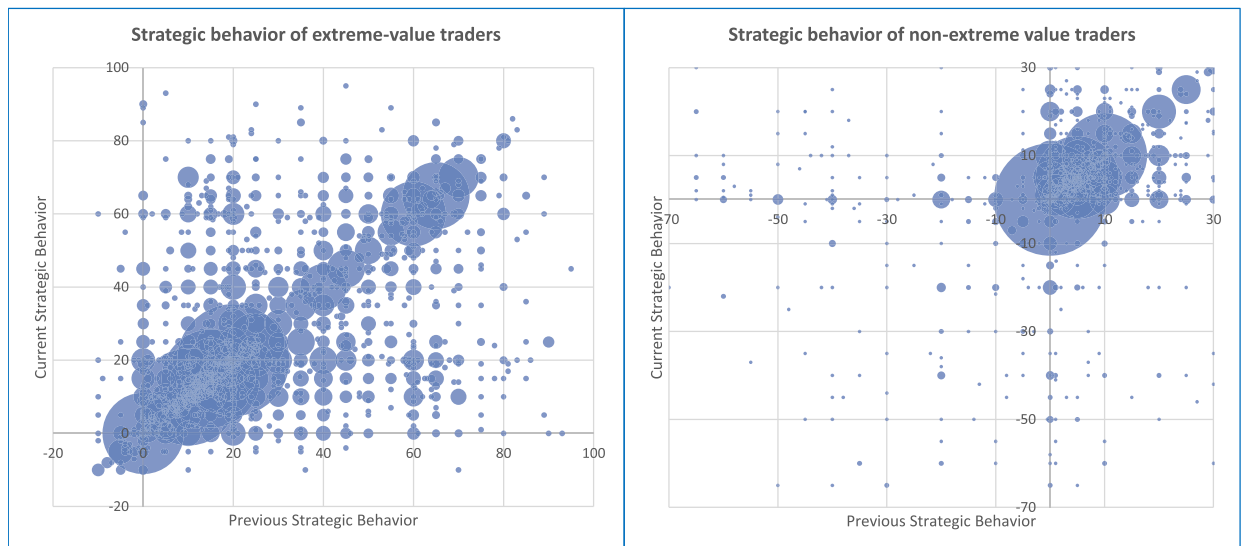
³⁴ Again, since our interest is in the treatment effects, not in prediction, linear regression coefficients approximate well the average treatment effects. Probit regressions, reported in Section S3.3 of Supplementary Materials online, lead to similar results. So do regressions on the combined sample and with standard errors clustered by individuals rather than matching groups, also reported in Section S3.3 of Supplementary Materials online.

³⁵ Non-parametric tests in Section S3.2 of Supplementary Materials online confirm the significance of this difference but do not find the differences between treatments S and N significant for either type of traders.

Table 5
Effects of treatment variations on naive behavior.

Variables	Dependent variable: Naive behavior (Naive)			
	Extreme		Non-extreme	
	I	II	III	IV
Treatment N	0.030 (0.019)	0.029 (0.019)	−0.014 (0.041)	−0.015 (0.040)
Frame DM	0.008 (0.020)	0.009 (0.020)	0.119*** (0.041)	0.120*** (0.041)
Setting “25”	0.002 (0.020)	0.002 (0.020)	0.004 (0.037)	0.005 (0.037)
Seller		−0.059*** (0.011)		0.080*** (0.022)
Period		< 0.001 (< 0.001)		0.002** (0.001)
Constant	0.051*** (0.013)	0.072*** (0.014)	0.243*** (0.048)	0.159*** (0.046)
Observations	5411		5469	
R ²	< 0.01		0.02	

Note: Naive is a dummy variable equal to 1 if strategic behavior is 0, and 0 otherwise. Extreme: values \underline{c} and \bar{v} ; non-extreme: values \underline{v} and \bar{c} . Treatment S, frame BA, and setting “20” are the base categories in all columns; Buyer is the base category in columns II and IV. Standard errors clustered by 32 matching groups in parentheses. *** - significant at 1% level; ** - significant at 5% level.



Note: Strategic behavior is $a - c$ or $v - b$. Current strategic behavior on the vertical axis; previous strategic behavior by the same subject (in the same role and with the same cost/value) on the horizontal axis. All data for which previous strategic behavior is available. Size of the circles corresponds to the number of observations. Left panel: extreme values (\bar{c} and \bar{v}); right panels: non-extreme values (\underline{v} and \bar{c}).

Fig. 3. Current and previous individual strategic behavior.

5. Influence of history on individual behavior

In the previous section, we saw that the average strategic behavior is not affected much by the treatment, frame, or setting variations. Moreover, the average behavior does not vary much over time. In this section, we focus on the time series of (disaggregated) individual behavior to argue that individual strategic behavior is different across subjects and appears to be influenced by history of play.

Fig. 3 shows how the current strategic behavior in individual observations (on the vertical axis) relates to the previous strategic behavior (Previous SB) by the same individual the last time he/she was in the same role and with the same cost/value (on the horizontal axis; depending on the assignment of roles and the realization of cost and values, it may be one or several periods ago). The left panel is for observations with extreme values and the right panel is for non-extreme values.

The figure shows that often the current behavior is the same as the previous one, but there is also variability around the 45° line, as subjects change their behavior. We formally check the association of the current behavior with the experienced history

Table 6
Effects of experienced history on behavior.

Dependent variable: Strategic behavior			Dependent variable: Naive		
Variables	Extreme	Non-extreme	Variables	Extreme	Non-extreme
	I	II		III	IV
Previous SB	0.660*** (0.017)	0.406*** (0.035)	Previous Naive	0.588*** (0.045)	0.607*** (0.026)
History-Best-Response	0.037*** (0.013)	0.063*** (0.017)	Payoff-Gain-History-BR	−0.0004 (0.001)	−0.003 (0.005)
Last-Best-Response	0.027** (0.010)	0.079*** (0.017)	Payoff-Gain-Last-BR	−0.0003* (0.0002)	−0.005*** (0.001)
Constant	6.096*** (0.748)	1.531*** (0.355)	Constant	0.040*** (0.007)	0.132*** (0.016)
Observations	4867	4925	Observations	4867	4925
R ²	0.44	0.19	R ²	0.34	0.37

Note: Strategic behavior is $a - c$ and $v - b$. Naive is a dummy variable equal 1 if strategic behavior is 0, and 0 otherwise. Extreme: values \underline{c} and \bar{v} ; non-extreme: values \bar{v} and \bar{c} . Previous SB is strategic behavior by the same subject last time the subject was in the same role and had the same cost/value. History-Best-Response is the strategic behavior at the best response to experienced history; Payoff-Gain-History-BR is the difference in payoff between this best response and naive play. Last-Best-Response is the strategic behavior at the best response to the last observation and Payoff-Gain-Last-BR is the difference in payoff between this best response and naive play. Standard errors clustered by 32 matching groups in parentheses. *** - significant at 1% level; ** - significant at 5% level; * - significant at 10% level.

by regressing current strategic behavior on previous strategic behavior, while also including variables measuring the effect of best response. Variable History-Best-Response is equal to how much strategic behavior should be if subject i were playing best response (BR) to the experienced history up to period t .³⁶ Similarly, variable Last-Best-Response is equal to the strategic behavior at best response to the last observation of the opponent's play by subject i in period t . These two variables measure the “optimal” strategic behavior in view of (long or short memory of) the player's history. Columns I–II of Table 6 report the regression results, separately for samples with extreme and non-extreme values. Columns III–IV of the table report regressions of the variable Naive on previous naive behavior (equal 1 if previous strategic behavior is 0, and 0 otherwise), and variables measuring incentives to deviate from naive play: variable Payoff-Gain-History-BR is the payoff difference between playing best response to the whole experienced history and playing naively,³⁷ and Payoff-Gain-Last-BR is the payoff difference between playing best response to the last observation and playing naively.

All the variables representing a subject's history in the game have significant coefficients in the regression reported in columns I–II. The extent of strategic behavior is correlated across time: the more strategic a subject was before, the greater the extent of the current strategic behavior by the subject. The coefficient on Previous SB is significantly smaller than 1 though; thus players do not always keep playing the same strategy. An increase in the distance between the best response to either the entire history of play or to the last observation and the player's value/cost increases the extent of strategic behavior too. Subjects who played (more) strategically tend to continue playing (more) strategically, but they also react to the history by changing towards best responses to their experience. From columns III–IV, playing naively is also correlated with playing naively in the past. A higher payoff difference between playing best response to past experience (whole history or last observation) and playing naively is associated with less naive play, although the effect is significant only for the payoff difference in the last observation.

We note the above findings about individual behavior in the next result:

Result 3. (Individual strategic behavior)

- The extent of strategic behavior in the past is positively correlated with the extent of strategic behavior now; strategic behavior also changes towards best responses to previous observations (both distant in the past and immediate).
- Playing naively in the past is positively correlated with playing naively now.

Another take from the previous section is the relatively large proportion of observations of naive play (also noticeable in Fig. 3). In the double auction, 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 while leading to a better price. In the N treatment, naive behavior is not consistent with any equilibrium since the buyer can obtain a higher expected payoff by bidding equal to the highest ask strictly below her value (which can be \bar{c} for the high-value buyer and \underline{c} for the low-value buyer). Similarly, the seller would prefer to set ask equal the lowest bid strictly above her cost (which can be \bar{v} for the high-cost seller

³⁶ The variable is constructed as follows. In the double auction game, the (seller's, for example) best response to observed bids is equal to either one of those bids or to the seller's own cost. For example, suppose that a seller with cost 10 observed bids 30, 70, 80, 30 in the previous periods. We calculate seller's payoff from asking 10, 30, 70, 80 against the observed history of bids; in this particular case, both 30 and 70 are best responses to that history. Such ties are broken in favor of the lowest strategic behavior; the value of History-Best-Response is thus $30 - 10 = 20$.

³⁷ This payoff difference is easy to calculate once we have determined the best response to any experienced history, as discussed in the previous footnote.

and \underline{v} for the low-cost seller). However, in the S treatment, naive behavior by traders with non-extreme values is consistent with a fully efficient equilibrium and thus not necessarily irrational. Naive play is not a best response given players' experience though: in the second part of the experiment (rounds 21–40), traders with non-extreme values play naively 29% of times, while in less than 5% of the times the best response to the observed history of play is within 0–4 distance from the cost/value.

One way to explain both the large proportion of naive play among non-extreme value traders and the lower such proportion among extreme-value traders is that subjects have (non-monetary) costs of lying (or, more generally, disutility from deviating from cost/value). 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 bid/ask is different from value/cost.³⁸ 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 setting bid equal z_1 , whereas the high-value buyer gains $\frac{1}{4}(z_2 - z_1) - x$. 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 sets bid different from value – which is consistent with a greater proportion of naive play by the traders with non-extreme values in the experiment. That the non-extreme value traders have a lower expected gain from lying than the 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.

Overall, past individual choices appear to influence present choices but the choices also react to the previous experience in the manner one can expect, towards best responses to observations. There is also some evidence that subjects may have disutility from deviation from cost/value. We use these results in the next section to motivate a dynamic model with reinforcement and noisy best response that can account for (some of) the properties of the data from the experiment.

6. Behavioral models of individual choices

Among the points learned from the previous sections are that the average behavior of the experiment participants and efficiency in the matches between human subjects are similar across treatments S and N. This zero indirect effect is somewhat surprising because, as discussed in Section 2, presence of naive traders dilutes the tradeoff between probability and terms of trade, which could incentivize more strategic behavior in treatment N. The average strategic behavior (for example, 25.54 for extreme and 4.26 for non-extreme value traders in setting “20”) is also incompatible with playing a pure-strategy equilibrium with a small noise, hinting at the possibility that a non-equilibrium model may have a better fit with data. Even though the average behavior and the efficiency level hardly change over time, individual behavior adjusts to some extent to the previous play. Finally, our findings suggest subjects experience disutility from deviating from cost/value. In this section, we introduce a dynamic model of behavior that reconciles these findings. We also compare this model to other models with a similar number of parameters, showing that it has a better fit to the data.

We build the models on the discrete set of bids/asks as multiples of 5 between 0 and 100 (the range used in the experiment): let $J = \{0, 5, \dots, 95, 100\}$. This set approximates the strategy set actually used: even though in our experiment bids and asks could be decimal numbers with up to two digits after the decimal point, 80% of choices were multiples of 5.³⁹

6.1. A dynamic model

Consider the uniform strategy on J of one trader. Suppose that the choices of the other trader follow a logistic distribution based on the expected payoffs against this uniform strategy, i.e.,

$$Prob(j) = \frac{e^{\lambda Eu(j)}}{\sum_{i \in J} e^{\lambda Eu(i)}}, \quad (1)$$

where $j \in J$ and $Eu(i)$ is the expected payoff of strategy i against the uniform strategy on J . Parameter λ measures the noise in choosing the best response: for $\lambda > 0$, strategies with higher payoffs have a higher probability to be chosen than strategies with lower payoffs; as $\lambda \rightarrow \infty$ the best response is chosen with probability approaching 1.⁴⁰

The previous paragraph describes the choice of a subject with level-1 strategic sophistication (see, for example, Stahl and Wilson, 1994). We assume that this models the initial choice of a strategy.⁴¹ To model the dynamics, we use the experience-weighted attraction (EWA) learning model (Camerer and Ho, 1999).⁴² In the EWA model, players form attractions for each strategy and then select a strategy according to the logistic equation (1) based on these attractions in place of expected payoffs. Attractions are updated based on experience, with those strategies that were played and those that either performed or would have performed best receiving more weight. The EWA model thus combines features of both reinforcement and belief-based models of learning. Strategies

³⁸ For example, Ellingsen and Johannesson (2004) and Serra-Garcia et al. (2013) consider models with such a fixed cost of lying.

³⁹ Note that this discretization preserves at least one fully-efficient equilibrium without naive traders in our experimental treatments: equilibrium strategies $Z_s^{str}(\underline{c}) = Z_b^{str}(\underline{v}) = 30$ are in J , as well as $Z_s^{str}(\bar{c}) = Z_b^{str}(\bar{v}) = 70$.

⁴⁰ This logistic choice is used, for example, 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 experience-weighted attraction model (EWA; Camerer and Ho, 1999).

⁴¹ Although other specifications of initial behavior are possible, level-1 strategic sophistication is a reasonable compromise reflecting both the complexity of the game and individual incentives to get a good payoff.

⁴² The EWA model has a better fit to observed behavior in a large class of games compared to both quantal response equilibrium and other dynamic models (Camerer and Ho, 1999; Ho et al., 2007) and has been used more recently to explain experimental data (e.g., Hong et al., 2021, and Masiliūnas, 2023).

that tended to do better in past are reinforced and, at the same time, behavior is updated based on how well a strategy would have performed against the other players' history of play. As discussed in the previous section (Result 3), these properties of the adjustment of strategies appear to be at least partially present in the data from our experiment.

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

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

where π_{it} is the payoff that strategy i obtained (or would have obtained)⁴³ 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 and the initial experience count N_0 are parameters.⁴⁴ While the initial attractions A_{i0} could also be treated as free parameters, as discussed above, we set them equal to the expected payoff of strategy i against the uniform distribution of opponent's choice on J , modeling level-1 strategic sophistication in initial choices.

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 by equation (2). For the current role, both sets of attractions (for extreme and non-extreme cost/value) are updated (since a player can use opponent's strategy to evaluate what would have happened if cost/value were different) but $I_{it} = 1$ only for the current cost/value (and actual strategy i played).

Motivated by the observed high frequency of naive behavior in the experiment, we consider that the game payoffs also include disutility of deviating from cost/value, with fixed component F (that applies only if ask/bid is different from cost/value) and marginal component f for each unit of deviation from cost/value.⁴⁵ Our model thus includes 7 parameters: 5 parameters (λ , δ , φ , ρ , N_0) of the EWA model and 2 parameters (F , f) of the disutility of deviating from cost/value.

We estimate the model by maximum likelihood on the data from 10 sessions with setting "20" (we use 6 sessions with setting "25" for validation of the estimates, as described below).⁴⁶ The estimated parameters are $\lambda = 0.231$, $\delta = 0.525$, $\varphi = 0.848$, $\rho = 0.820$, $N_0 = 4.692$ for the EWA model parameters⁴⁷ and $F = 4.076$, $f = 0.081$ for the disutility parameters, with the log-likelihood value $LL = -11167.0$.

We use these values of the parameters to simulate experimental sessions. Each simulated run has a matching group of 8 subjects, randomly matched for 40 rounds and each using the model above to make decisions. The average results of 100 simulated runs for treatments S and N are presented in Figs. 4 and 5.⁴⁸ 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 in simulations is noticeably lower than the expected efficiency in treatment N, as in the experiment.⁴⁹ There is little difference in the extent of strategic behavior between N and S treatments in simulations for both extreme and non-extreme value traders. The simulated averages are also very close to those from the data – to the extent that some thin lines representing the experimental data are hardly visible, being covered by the thicker lines of the simulated data. Even though the model is dynamic, there is little change in either efficiency or average strategic behavior over time.

The parameters used in the simulations are estimated only on the data from setting "20". The right panels of the figures provide a visual validation that the model with these parameters also fits setting "25": the averages of the simulated data are also close to the averages of the experimental data in these panels. As a more formal validation measure we use the mean absolute deviations between the simulated and experimental averages of strategic behavior for blocks of 5 periods used to draw Fig. 5. For setting "20" in the left panels, this measure is 1.260; for setting "25" in the right panels, the measure is 1.349. That the measure is only 0.089 higher out-of-sample (while, as can be seen in the figure, the averages of strategic behavior vary on the scale between 0 and 30) shows that the model also works reasonably well for setting "25", indicating little overfitting.⁵⁰

⁴³ If payoffs in a game can be negative, then the interpretation of the term $(\delta + (1 - \delta) I_{it}) \pi_{it}$ as reinforcement is less clear (the multiplication of a negative payoff by a number less than 1 increases it, while such multiplication for a positive payoff decreases it). We therefore add 70 to all payoffs in the game to make them all non-negative. If initial attractions are based on expected payoffs, such an additive transformation does not change the choice probabilities in equation (1) for round 1. For subsequent rounds it adds inertia in the choice of strategies by increasing the reinforcement of the chosen strategy (Xie, 2021).

⁴⁴ Parameter δ reflects the mix between reinforcement and belief-based learning. If $\delta = 0$, then only the previously played strategy is reinforced, as in reinforcement learning, whereas if $\delta = 1$, then all strategies are considered for updating, as in belief-based learning. Parameters φ , ρ , and N_0 are related to the speed of learning, measuring the decay of previous attractions, the decay of experience, and the strength of the prior, respectively. See Camerer and Ho (1999) for further explanation.

⁴⁵ This form of deviation cost combines the forms of lying costs discussed, e.g., in Ellingsen and Johannesson (2004), Kartik (2009), Serra-Garcia et al. (2013) and Gneezy et al. (2018). Masiliūnas (2023) also considers payoff modifications in EWA learning.


⁴⁶ Details of the estimation procedure can be found in Section S4.1 of Supplementary Materials online.

⁴⁷ The estimates are consistent with other estimates in the literature. For example, Stahl and Wilson (1994), McKelvey and Palfrey (1995), Camerer and Ho (1999), and, more recently, Nunnari and Zapal (2016), Hong et al. (2021), Masiliūnas (2023), estimate λ in range [0.1, 0.5] (adjusted for our range of payoffs). Estimates for δ , φ and ρ are not far from those discussed in Camerer and Ho (1999). That δ is not close to 0 or 1 indicates that both reinforcement and best-responding are important in the model.

⁴⁸ 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.

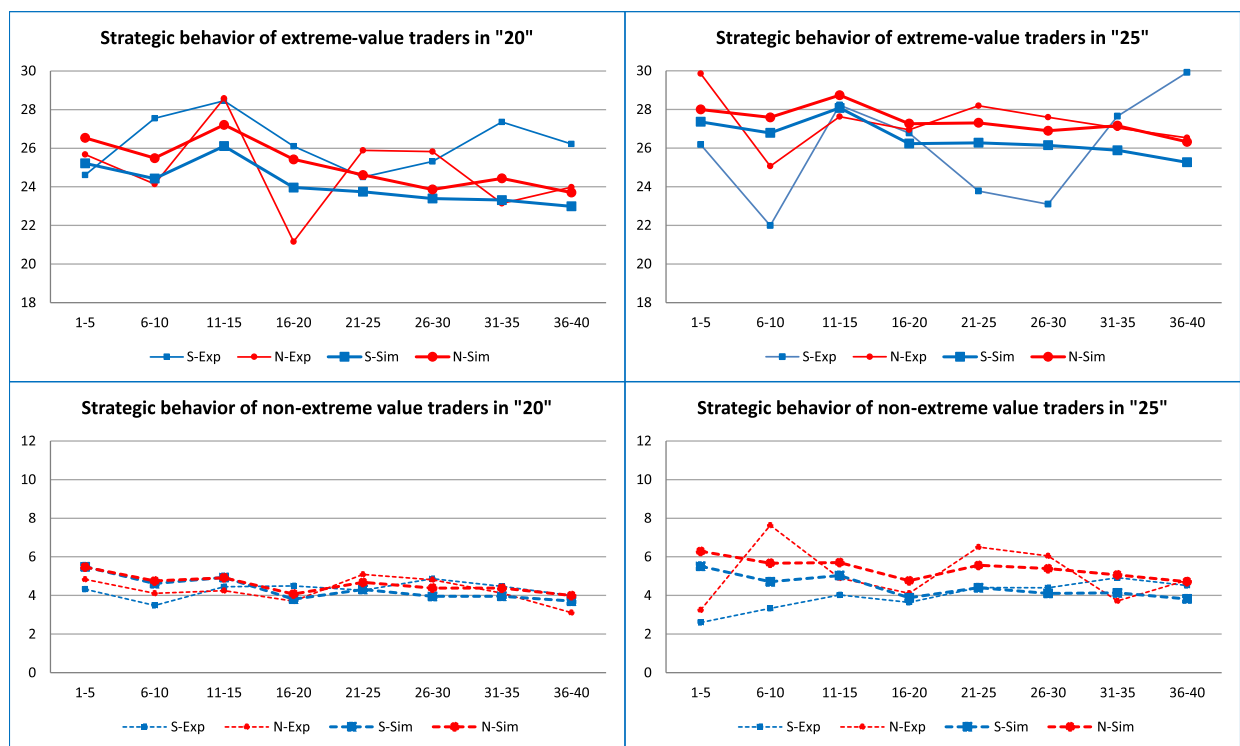
⁴⁹ The human matches efficiency in treatment N in simulations (not shown to make the figure less cluttered) is only slightly lower than the realized efficiency in treatment S.

⁵⁰ As we discuss in the next subsection, the EWA model also has a better fit than other models, both in- and out-of-sample.



Note: Blocks of periods on the horizontal axis; fraction of surplus realized on the vertical axis (expected efficiency). S: treatment S; N: treatment N. -Exp (thinner lines): experimental data; -Sim (thicker lines): simulated data.

Fig. 4. Efficiency in simulations and in actual sessions.



Note: Blocks of periods on the horizontal axis; average strategic behavior ($v - b$ and $a - c$) on the vertical axis. S: treatment S; N: treatment N. -Exp (thinner lines): experimental data; -Sim (thicker lines): simulated data. Top panels: extreme values (\underline{c} and \bar{v}); bottom panels: non-extreme values (\underline{v} and \bar{c}).

Fig. 5. Strategic behavior in simulations and in actual sessions.

6.2. Comparison with other models

Given that average strategic behavior changes little over time in the experiment, alternative models of behavior can assume that subjects use the same strategy in all periods.⁵¹

We first check whether a noisy equilibrium-like behavior can explain our data. For this, we look at the actual distributions of behavior in each matching group.⁵² If behavior is close to an equilibrium, it should fit with (possibly noisy, as represented by

⁵¹ Details of the estimation procedures for the two static models below can be found in Section S4.2 of Supplementary Materials online.

⁵² Since the double auction game has multiple equilibria, different matching groups can in principle be close to different equilibria.

Table 7
Comparison of models in- and out-of-sample.

Model	<i>k</i>	In-sample (setting “20”)			Out-of-sample (setting “25”)		
		LL	AIC	BIC	LL	AIC	BIC
EWA (deviation cost)	7	−11167.0	22348.0	22395.8	−7119.8	14253.6	14301.3
EWA (fairness)	8	−11347.2	22710.4	22765.0	−7210.3	14436.7	14491.3
“Equilibrium-like”	3	−15502.1	31010.1	31030.6	−9972.3	19950.6	19971.1
“Mark-up”	12	−14526.1	29076.1	29158.0	−8947.8	17919.7	18001.6

Note: *k* is the number of parameters in a model. *n* = 6800 observations for setting “20”, *n* = 4080 for setting “25”. LL: log-likelihood on the parameters estimated for setting “20”; AIC: Akaike Information Criterion ($2k - 2LL$); BIC: Bayesian Information Criterion ($\ln(n)k - 2LL$). Models are as described in text.

equation (1)) best response to these distributions. We estimate such a model by maximum likelihood with 3 parameters: best response noise λ and two parameters F, f allowing for disutility from deviating from value/cost. The estimated maximum log-likelihood value is $LL = -15502.1$, much lower than the value for the EWA model.

Constant average strategic behavior is also consistent with a model in which subjects use a static strategy of setting ask somewhat above cost (or bid somewhat below value). To allow for the observed persistence of individual behavior over time, we estimate such a model where behavior of a subject, for a given type (extreme or non-extreme), comes from one of two distributions over J (or from a mixture of these distributions). Each distribution is based on a normal distribution with two parameters: mean strategy (modeling average mark-up) and variance (modeling possible deviations from the average mark-up).⁵³ Using two distributions allows for subjects who may be naive, strategic, or a mixture of the two. The maximized log-likelihood value of this model is $LL = -14526.1$, again worse than the value for the EWA model.

Thus, even though the average behavior over the 40 rounds of experiment changes little, the dynamic strategy adjustment model fits the behavior better than the two static models. While it is possible that some subjects do not change their behavior over time, sufficiently many subjects modify their choices, justifying the consideration of the dynamic model.

In bilateral bargaining situations, players can also have preferences for a fair distribution of the trade surplus, with utility depending not only on player's own payoff but also on how different it is from the other player's payoff.⁵⁴ In our setting, since a trader does not know the cost or the valuation of the other trader, what is fair ex ante is not easy to determine but a player can evaluate the fairness of an outcome ex post, if cost and value were observed or can be inferred from ask/bid.⁵⁵

We estimate a mixture EWA model, allowing a proportion of subjects to have fairness preferences as in Fehr and Schmidt (1999). The maximized log-likelihood of this model is $LL = -11347.2$, worse than the model with disutility of deviations from cost/value, even though the mixture model has one more parameter.⁵⁶ In our experiment, the behavior of extreme-value traders can be consistent with inequity aversion, but many non-extreme traders make asks/bids close to their cost/value, which is difficult to reconcile with fairness, since it leaves a large part of the surplus to the other trader. While it is possible that some subjects have fairness preferences, to fit the data such preferences would need to differ depending on the position in the game (extreme or non-extreme). The model with disutility of deviating from cost/value, with parameters that do not depend on the position in the game, on the other hand, provides a more parsimonious explanation of behavior.

Table 7 summarizes the comparison of the models. Both the raw log-likelihood and the values of information criteria are better for the EWA model with disutility of deviating from cost/value than for the alternative models. This holds for the sample used for estimation (data for setting “20”), as well as out-of-sample (data for setting “25”),⁵⁷ again confirming that the model, despite its large number of parameters, is not overfitted in-sample.

The comparison shows that the EWA learning model has the best fit among the considered models. There are still some differences between the data from EWA simulations and from experiment. For example, the behavior of non-extreme-value traders is more dispersed in the simulations than in the experiment, with such traders more often bidding above value or asking below cost in simulations than in the experiment. The simulations have a built-in symmetry across roles and do not distinguish between the BA and DM frames.⁵⁸ Nevertheless, the simulations demonstrate that a learning model can explain the behavior in the experiments to a large extent. In particular, the simulations show that with reinforcement of strategies and learning noisy best response, convergence to the efficient equilibrium is difficult, and that the presence of naive traders does not necessarily lead to noticeably more aggressive behavior and less efficiency, confirming what we observe in our experiment.

⁵³ These normal distributions are adjusted to probability distributions on the discrete set J , as described in Section S4.2 of Supplementary Materials online.

⁵⁴ Roth (1995) discusses fairness in bargaining experiments; Fehr and Schmidt (1999) and Bolton and Ockenfels (2000) are among the first to model such preferences.

⁵⁵ Section S4.3 of Supplementary Materials online describes in more details how we model these preferences in our setting.

⁵⁶ The details of estimation are presented in Section S4.3 of Supplementary Materials online.

⁵⁷ For the “mark-up” model the parameters estimated from setting “20” are modified for setting “25” by adjusting the mean strategy of the extreme value traders, as described in Section S4.2 of Supplementary Materials online.

⁵⁸ A possible way to get simulations have a better fit and reproduce observed differences across roles and frames is to have heterogeneity across subjects, with some having higher disutility from deviating from value/cost or some never bidding above value or below cost.

7. 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, who set bid/ask equal to value/cost, this efficient equilibrium disappears; thus, the presence of naive traders may decrease efficiency.

We ran an experiment to see whether the presence of such naive traders would reduce efficiency but did not find this effect. We found that the experimental subjects did not play more strategically after an introduction of artificial naive traders, and did not play according to the efficient equilibrium in the treatment without such naive traders. Instead, subjects tended to reinforce strategies that worked well in the past as well as adjust behavior toward strategies that would have done better against the history of other player's play. Motivated by this finding, we presented a learning model, with initial choices described by noisy best responses according to low levels of strategic sophistication and subsequent adjustment based on past experience, that replicates average measures of efficiency and behavior from the experimental data (while also providing a better fit than other models with a comparable number of parameters).

Taken at face value, our finding that play may be better described by learning starting from low levels of strategic sophistication appears to contradict the previous literature (Radner and Schotter, 1989; Valley et al., 2002; McGinn et al., 2003; Ellingsen et al., 2009). That literature found that the constrained-efficient linear equilibrium of the double auction describes the data well even though the information setting in those experiments (uniform distributions of value and cost on an interval) is more complex than our setting of two-point value/cost distributions.⁵⁹ There is one way to resolve these seemingly contradictory findings. In the model with identical uniform distributions of values, observed behavior can be consistent with the linear equilibrium of the double auction even though players have low levels of strategic sophistication and do not play the equilibrium strategies. For instance, assuming uniform-random play at level-0, expected behavior is equal to the linear equilibrium strategy if there are 25% of (non-noisy) level-1 players and 75% of (non-noisy) level-2 players.⁶⁰ Thus, it is possible that players' behavior in that setting appears close to linear equilibrium although it is better described by non-equilibrium models like level- k and learning dynamics – of course, a rigorous reanalysis of data from previous experiments is necessary to test this conjecture.

The conclusions reached in this paper, which contrast some of the previous literature, call for further 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, and whether and how its extent changes when parameters of the situation change), as well as the possibility (or impossibility) of achieving full efficiency in practice, also (and perhaps especially) when it is theoretically feasible in equilibrium.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Appendix A. Proofs

Proposition 1. 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

$$\max\{A(\underline{c}, \underline{v}), B(z_2, \underline{c}, \underline{v}), C(z_2, \underline{c}, \underline{v}, \bar{v})\} \leq z_1 \leq D(\underline{c}, \underline{v}) \quad (3)$$

and

$$E(\bar{c}, \bar{v}) \leq z_2 \leq \min\{F(\bar{c}, \bar{v}), G(z_1, \bar{c}, \bar{v}), H(z_1, \bar{c}, \bar{v}, \underline{c})\}. \quad (4)$$

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. Similarly, the bid of a strategic buyer cannot be strictly above the value. In a fully efficient equilibrium, $Z_s^{str}(\underline{c}) \leq Z_b^{str}(\underline{v})$ and $Z_s^{str}(\bar{c}) \leq Z_b^{str}(\bar{v})$. Thus

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

⁵⁹ Valley et al. (2002) and McGinn et al. (2003) find that pre-play communication generates efficiency gains that are greater than the theoretical maximum in the setting with uniform distributions. In Possajennikov and Saran (2022), we show that allowing pre-play communication in our setting does not always lead to full efficiency.

⁶⁰ For example, buyer's level-1 and level-2 strategies are $2/3v$ and $2/3v + 1/9$, respectively. If 25% of buyers are level-1 and the rest level-2, then the buyer's expected behavior equals $2/3v + 1/12$, which is the linear equilibrium strategy of the buyer in the game. See Crawford (2021) for further details of level-1 and level-2 strategies in this setting.

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

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

$$(1 - \varepsilon) \left(\frac{1}{2}(z_1 - \underline{c}) + \frac{1}{2} \left(\frac{z_1 + z_2}{2} - \underline{c} \right) \right) + \varepsilon \left(\frac{1}{2} \left(\frac{z_1 + \underline{v}}{2} - \underline{c} \right) + \frac{1}{2} \left(\frac{z_1 + \bar{v}}{2} - \underline{c} \right) \right) \quad (5)$$

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

$$(1 - \varepsilon) \left(\frac{1}{2} \cdot 0 + \frac{1}{2} \left(\frac{\underline{v} + z_2}{2} - \underline{c} \right) \right) + \varepsilon \left(\frac{1}{2} \left(\frac{\underline{v} + \underline{v}}{2} - \underline{c} \right) + \frac{1}{2} \left(\frac{\underline{v} + \bar{v}}{2} - \underline{c} \right) \right). \quad (6)$$

The difference between payoffs (5) and (6) is non-negative if and only if $z_1 \geq A(\underline{c}, \underline{v})$. The seller would get the expected payoff

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

from setting ask equal z_2 . The difference between payoffs (5) and (7) 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

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

The difference between payoffs (5) and (8) is non-negative if and only if $z_1 \geq C(z_2, \underline{c}, \underline{v}, \bar{v})$. All other deviations give 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

$$(1 - \varepsilon) \left(\frac{1}{2}(\underline{v} - z_1) + \frac{1}{2} \cdot 0 \right) + \varepsilon \left(\frac{1}{2} \left(\underline{v} - \frac{z_1 + \underline{c}}{2} \right) + \frac{1}{2} \cdot 0 \right) \quad (9)$$

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

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

The difference in payoffs (9) and (10) 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 .

If $\varepsilon = 0$, then $A(\underline{c}, \underline{v}) = \frac{1}{3}(2\underline{c} + \underline{v})$, $B(z_2, \underline{c}, \underline{v}) = \frac{1}{3}(z_2 + 2\underline{c})$ and $C(z_2, \underline{c}, \underline{v}, \bar{v}) = \frac{1}{3}(4\underline{c} - z_2)$. For $\bar{c} \leq z_2 \leq \bar{v}$, $\max\{A(\underline{c}, \underline{v}), B(z_2, \underline{c}, \underline{v}), C(z_2, \underline{c}, \underline{v}, \bar{v})\} = \frac{1}{3}(z_2 + 2\underline{c})$. Also, $D(\underline{c}, \underline{v}) = \underline{v}$. Therefore one inequality for fully efficient equilibrium if $\varepsilon = 0$ is $\frac{2}{3}\underline{c} + \frac{1}{3}z_2 \leq z_1 \leq \underline{v}$. Analogously, substituting $\varepsilon = 0$ into $E(\bar{c}, \bar{v})$, $F(\bar{c}, \bar{v})$, $G(z_1, \bar{c}, \bar{v})$ and $H(z_1, \bar{c}, \bar{v}, \underline{c})$ gives $\bar{c} \leq z_2 \leq \frac{2}{3}\bar{v} + \frac{1}{3}z_1$. \square

Proposition 2. 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, if $\varepsilon = 0$, 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 (3) in Proposition 1 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 (4), $E(70, 90) \leq z_2$, 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 1. \square

Proposition 3. 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 (3) in Proposition 1 means that

$$z_1 \leq D(\underline{c}, \underline{v}) = \frac{1}{2-\varepsilon}(60 - 55\varepsilon). \quad (11)$$

Since $G(z_1, \bar{c}, \bar{v}) = (z_1 + 190 - 70\epsilon)/(3 - \epsilon)$, inequality (4), combined with (11) implies $(3 - \epsilon)(2 - \epsilon)z_2 \leq 440 - 385\epsilon + 70\epsilon^2$. Since $E(\bar{c}, \bar{v}) = (140 - 45\epsilon)/(2 - \epsilon)$, inequality (4) also implies $140 - 45\epsilon \leq (2 - \epsilon)z_2$. Thus $(3 - \epsilon)(140 - 45\epsilon) \leq 440 - 385\epsilon + 70\epsilon^2$, which is equivalent to $0 \leq 4 - 22\epsilon + 5\epsilon^2$. This holds if $\epsilon \leq \bar{\epsilon} = (11 - \sqrt{101})/5 \approx 0.190$ and does not hold for $\epsilon > \bar{\epsilon}$. Therefore for $\epsilon > \bar{\epsilon}$ there is no z_1, z_2 that satisfy the inequalities of Proposition 1.

Take

$$z_1 = D(5, 30) = \frac{1}{2 - \epsilon}(60 - 55\epsilon) \quad \text{and} \quad z_2 = E(70, 95) = \frac{1}{2 - \epsilon}(140 - 45\epsilon). \quad (12)$$

By construction, $z_1 \leq D(5, 30)$. It holds that $A(5, 30) \leq z_1$ if $(2 - \epsilon)(40 + 20\epsilon) \leq (3 - \epsilon)(60 - 55\epsilon)$. This holds if $\epsilon \leq \sqrt{129}/6 - 9/6 \approx 0.393$. For the given z_2 , $B(z_2, 5, 30) = (160 - 115\epsilon + 30\epsilon^2)/((2 - \epsilon)(3 - \epsilon))$. Then $B(z_2, 5, 30) \leq z_1$ if $32 - 23\epsilon + 6\epsilon^2 \leq (3 - \epsilon)(12 - 11\epsilon)$, or $\epsilon \leq (11 - \sqrt{101})/5 = \bar{\epsilon}$. Finally, $C(z_2, 5, 30, 95) = (-100 + 275\epsilon - 100\epsilon^2)/((2 - \epsilon)(3 - \epsilon)) \leq z_1$ if $-20 + 55\epsilon - 20\epsilon^2 \leq (3 - \epsilon)(12 - 11\epsilon)$. This holds if $\epsilon \leq 50/31 - 2\sqrt{191}/31 \approx 0.721$. Therefore for z_1 and z_2 in (12), inequalities in (3) of Proposition 1 are satisfied if $\epsilon \leq \bar{\epsilon}$. Analogous reasoning shows that inequalities in (4) are also satisfied. Thus z_1 and z_2 in (12) constitute a fully efficient equilibrium for any $\epsilon \leq \bar{\epsilon}$. \square

Appendix B. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jebo.2023.10.003>.

References

- Abeler, J., Nosenzo, D., Raymond, C., 2019. Preferences for truth-telling. *Econometrica* 87, 1115–1153.
- Bolton, G.E., Ockenfels, A., 2000. ERC: a theory of equity, reciprocity, and competition. *Am. Econ. Rev.* 90, 166–193.
- 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. *Rev. Econ. Stud.* 54, 175–192.
- Chatterjee, K., Samuelson, W., 1983. Bargaining under incomplete information. *Oper. Res.* 31, 835–851.
- Crawford, V.P., 2021. Efficient mechanisms for level- k bilateral trading. *Games Econ. Behav.* 127, 80–101.
- Crawford, V.P., Costa-Gomes, M.A., Iriberry, N., 2013. Structural models of nonequilibrium strategic thinking: theory, evidence, and applications. *J. Econ. Lit.* 51, 5–62.
- Ellingsen, T., Johannesson, M., 2004. Promises, threats and fairness. *Econ. J.* 114, 397–420.
- Ellingsen, T., Johannesson, M., Lilja, J., Zetterqvist, H., 2009. Trust and truth. *Econ. J.* 119, 252–276.
- Fehr, E., Schmidt, K.M., 1999. A theory of fairness, competition, and cooperation. *Q. J. Econ.* 114, 817–868.
- Feri, F., Gantner, A., 2011. Bargaining or searching for a better price? – An experimental study. *Games Econ. Behav.* 72, 376–399.
- Fischbacher, U., 2007. z-Tree: Zurich toolbox for ready-made economic experiments. *Exp. Econ.* 10, 171–178.
- Fischbacher, U., Föllmi-Heusi, F., 2013. Lies in disguise: an experimental study on cheating. *J. Eur. Econ. Assoc.* 11, 525–547.
- Gneezy, U., 2005. Deception: the role of consequences. *Am. Econ. Rev.* 95, 384–394.
- Gneezy, U., Kajackaite, A., Sobel, J., 2018. Lying aversion and the size of the lie. *Am. Econ. Rev.* 108, 419–453.
- Greiner, B., 2015. Subject pool recruitment procedures: organizing experiments with ORSEE. *J. Econ. Sci. Assoc.* 1, 114–125.
- Ho, T.H., Camerer, C.F., Chong, J.-K., 2007. Self-tuning experience weighted attraction learning in games. *J. Econ. Theory* 133, 177–198.
- Hong, J., Moiras, S., Pouget, S., 2021. Learning in speculative bubbles: theory and experiment. *J. Econ. Behav. Organ.* 185, 1–26.
- Kartik, N., 2009. Strategic communication with lying costs. *Rev. Econ. Stud.* 76, 1359–1395.
- Kneeland, T., 2022. Mechanism design with level- k types: theory and an application to bilateral trade. *J. Econ. Theory* 201, 105421.
- Leininger, W., Linhart, P., Radner, R., 1989. Equilibria of the sealed-bid mechanism for bargaining with incomplete information. *J. Econ. Theory* 48, 63–106.
- Lundquist, T., Ellingsen, T., Gribbe, E., Johannesson, M., 2009. The aversion to lying. *J. Econ. Behav. Organ.* 70, 81–92.
- Masilūnas, A., 2023. Learning in rent-seeking contests with payoff risk and foregone payoff information. *Games Econ. Behav.* 140, 50–72.
- McGinn, K., Thompson, L., Bazerman, M., 2003. Dyadic processes of disclosure and reciprocity in bargaining with communication. *J. Behav. Decis. Mak.* 16, 17–34.
- McKelvey, R.D., Palfrey, T.R., 1995. Quantal response equilibria for normal form games. *Games Econ. Behav.* 10, 6–38.
- Myerson, R., Satterthwaite, M., 1983. Efficient mechanisms for bilateral trading. *J. Econ. Theory* 29, 265–281.
- Nunnari, S., Zapal, J., 2016. Gambler's fallacy and imperfect best response in legislative bargaining. *Games Econ. Behav.* 99, 275–294.
- Possajennikov, A., Saran, R., 2022. Private Value Bargaining with Communication: Does Cheap-Talk Lead to Efficiency? Working Paper.
- Radner, R., Schotter, A., 1989. The sealed-bid mechanism: an experimental study. *J. Econ. Theory* 48, 179–220.
- Roth, A.E., 1995. Bargaining experiments, Ch. 4. In: Kagel, J.H., Roth, A.E. (Eds.), *Handbook of Experimental Economics*. Princeton University Press, Princeton, pp. 253–345.
- Saran, R., 2011. Bilateral trading with naive traders. *Games Econ. Behav.* 72, 544–557.
- Saran, R., 2012. How Naiveté improves efficiency in trading with pre-play communication. *Econ. Lett.* 117 (1), 311–314.
- Satterthwaite, M., Williams, S., 1989. Bilateral trade with the sealed bid k -double auction: existence and efficiency. *J. Econ. Theory* 48, 107–133.
- Serra-Garcia, M., van Damme, E., Potters, J., 2013. Lying about what you know or about what you do? *J. Eur. Econ. Assoc.* 11, 1204–1229.
- 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. *J. Econ. Behav. Organ.* 25, 309–327.
- Valley, K., Thompson, L., Gibbons, R., Bazerman, M.H., 2002. How communication improves efficiency in bargaining games. *Games Econ. Behav.* 38, 127–155.
- Xie, E., 2021. Empirical properties and identification of adaptive learning models in behavioral game theory. *J. Econ. Behav. Organ.* 191, 798–821.