

Learning From Unrealized versus Realized Prices

M. Kathleen Ngangoue (New York University)
Georg Weizsäcker (Humboldt University Berlin)

Discussion Paper No. 66

January 18, 2018

Learning from unrealized versus realized prices*

M. Kathleen Ngangoué Georg Weizsäcker

January 2018

Abstract

Our experiments investigate the extent to which traders learn from the price, differentiating between situations where orders are submitted before versus after the price has realized. In simultaneous markets with bids that are conditional on the price, traders neglect the information conveyed by the hypothetical value of the price. In sequential markets where the price is known prior to the bid submission, traders react to price to an extent that is roughly consistent with the benchmark theory. The difference's robustness to a number of variations provides insights about the drivers of this effect (*JEL* D82, D81, C91).

1 Introduction

Market prices reflect much information about fundamental values. The extent to which traders are able to utilize this information has important welfare consequences but is difficult to measure as one often lacks control of the traders' restrictions, beliefs and preferences. One possibility to detect a bias in price inference is to modify the informational environment in a way that is irrelevant for rational traders. If trading reacts to a framing variation that is uninformative under rational expectations, the latter assumption is questionable. We focus on an important dimension of variability between markets, the conditionality of price. In *simultaneous* markets, the price realization is unknown to the traders at the time when they make their decisions—examples are financial markets with limit orders or other supply/demand function regimes. Theoretically, traders would incorporate the information of each possible price into their bids, as in the Rational Expectations Equilibrium prediction by Grossman (1976), *inter alia*.

*We benefited from helpful discussions with Dan Benjamin, Erik Eyster, Steffen Huck, Ryan Oprea, Emanuel Vespa, Roberto Weber, as well as with the audiences at UC Berkeley, UCL, DIW Berlin, UC San Diego, UC Santa Barbara, USC, FU Berlin, Hamburg University, the Berlin Behavioral Economics Workshops, and the conferences CED, EEA Meeting, ESWC, ICGT Stony Brook, SITE, TARK and THEEM. We also thank colleagues at Technical University Berlin and WZB for support in the conduct of the experiment. Financial support by the European Research Council (through Starting Grant 263412) and the German Science Foundation (through Collaborative Research Center TRR 190) is gratefully acknowledged. Ngangoué: New York University, kn44@nyu.edu. Weizsäcker: Humboldt-Universität zu Berlin and DIW Berlin, weizsaecker@hu-berlin.de

But the price information is hypothetical and traders may find it hard to make the correct inference in hypothetical conditions. A host of evidence on Winner’s Curse and other economic decision biases is consistent with this conjecture, as is the psychological evidence on accessibility (Kahneman, 2003) and contingent thinking (Evans, 2007).¹ Simultaneous asset markets with price-taking agents are a relevant point in case for such failures of contingent thinking; one that has not previously been researched, to our knowledge. In contrast, *sequential* markets—e.g. many quote-based markets and sequential auctions—have the traders know the price at which they can complete their trades. Here, it may still be nontrivial to learn from the price; but both the psychological research on contingent reasoning and the related economic experiments that include treatment variations where simultaneity is switched on and off (Carrillo and Palfrey 2011, Esponda and Vespa 2014; 2016 and Li 2016) suggest that the task is more accessible in a sequential trading mechanism than in a simultaneous one. Our series of experiments confirms this hypothesis, in a simple and non-strategic market environment where agents act as price takers.² In such an environment, the failure to learn from the price is especially noteworthy because the price explicitly reflects the asset value, conditional on the available information. To shed further light on the importance of this failure, we study its potential sources and discuss possible implications in financial markets.

The comparison between the two extreme trading mechanisms enables us to identify sets of trades that can be directly attributed to imperfect contingent thinking. We prefer avoiding claims about external validity but we note that the necessity to think contingently is ubiquitous in real-world markets, at various levels, despite the fact that a clear distinction between pure simultaneous and sequential markets vanishes. Order-driven markets, especially in the form of call auctions, require investors to supply liquidity without knowledge of the liquidity demand (Malinova and Park, 2013; Comerton-Forde et al., 2016). Examples of pure order-driven markets are the stock exchanges in Hong Kong, Japan and several other Asian countries, whereas the London SEAQ, for instance, functions as a pure quote-driven market.

Markets that represent hybrid versions of order- and quote-driven mechanisms also exhibit important features of simultaneous trading. For example, equity markets with low liquidity may be cleared throughout the day with periodically conducted call auctions; other markets open or close the day’s trading via call auctions. Additionally, an increasing flow of retail orders is internalized (Comerton-Forde et al., 2016). These orders are not executed on public exchanges but are executed internally through dark avenues or routed to different exchanges, making it difficult for retail investors to monitor the market conditions prior to trade. Thus, even for continuously traded assets the increasing market fragmentation and the increasing speed of trades force (slow) retail investors to post orders without precise knowledge of transaction prices, requiring contingent thinking.

¹Experiments analyzing the Winner’s Curse include, for example, Bazerman and Samuelson (1983); Kagel and Levin (1986); Kagel, Levin, Battalio, and Meyer (1989). For a thorough review on the Winner’s Curse literature see Kagel and Levin (2009).

²While technically incompatible, our evidence may be viewed as supporting the main idea of Li’s (2016) obvious strategy proofness: in a sequential market, the set of prices that are still possible is smaller than in simultaneous markets, enabling the trader to identify an optimal strategy.

The difference in informational efficiency between simultaneous and sequential trading mechanisms has been discussed both theoretically (e.g. Kyle 1985; Madhavan 1992; Pagano and Roell 1996) and experimentally (Schnitzlein, 1996; Theissen, 2000; Pouget, 2007). A consensus is that, in the presence of perfectly informed insiders, the temporal consolidation of orders in call auctions allows markets to aggregate information as efficiently as with continuous trading.³ With heterogeneous information, in contrast, the possibility to learn from market prices becomes essential when private information is at odds with the aggregate information, and determines the speed of price discovery. This holds in particular when new information flows into markets. Yet, an established pattern is that prices in real and experimental call markets adjust relatively slowly to incoming information (Amihud et al., 1997; Theissen, 2000). Contributing to a possible explanation of this pattern, we further document and examine the discrepancies between stylized simultaneous and sequential markets, with a focus on the extent to which traders learn from the price.

Our participants trade a single, risky, common-value asset. To trade optimally, a participant considers two pieces of information: her private signal and the information conveyed by the asset price. The latter is informative because it is influenced by the trading activity of another market participant who has additional information about the asset value. To manipulate the accessibility of the price information, we perform the experiment in two main treatments, simultaneous (SIM) versus sequential (SEQ). In treatment SIM, participants receive a private signal and submit a limit order. If the market price realizes below the limit, the trader buys one unit of the asset, otherwise she sells one unit.⁴ Despite the fact that the price has not yet realized, SIM traders would optimally infer the extent to which a high price indicates a high value and, thus, soften the demand's downward reaction to a higher price, relative to the case that the price is uninformative. The possibility that traders may fail to learn from hypothetical prices is examined by comparing to the treatment with sequential markets, SEQ, where the price is known when traders choose to buy or sell. Conditional thinking is not necessary here but treatments SIM and SEQ are nevertheless equivalent: they have isomorphic strategy sets and isomorphic mappings from strategies to payoffs.

Section 2 presents the experimental design in detail and Section 3 discusses our behavioral hypotheses. We present three benchmark predictions for comparison with the data: first, full naiveté, where the trader learns nothing from the price; second, the Bayes-Nash prediction, where a trader assumes that previous trades are fully rational and accounts for it; and third, the empirical best response that takes into account the actual distribution of previous trades, which may deviate from optimality. We use the latter as our main benchmark for optimality as it maximizes the traders' expected payments. That is, we ask whether naiveté fits the data better than the empirical best response, separately by treatment.

³Pouget's experimental call market is informationally efficient because of the high share of insiders, but liquidity provision in call markets deviates more from equilibrium predictions. This finding is consistent with ours and Pouget, too, assigns the deviation from equilibrium strategies to bounded rationality and partly to strategic uncertainty.

⁴Traders also have the option to reverse their limit order, selling at low prices and buying at high prices. This ensures the equivalence between the treatments, see Section 2. In each treatment, we restrict the trades to a single unit of supply or demand per trader.

The data analysis of Section 4 shows that the participants’ inference of information from the price varies substantially between simultaneous versus sequential markets. In SIM, participants often follow the prediction of the naive model, thus showing ignorance of the information contained in the price. Price matters mainly in its direct influence on the utility from trade—a buyer pays the price, a seller receives it. In contrast, in SEQ, where transaction prices are known beforehand, asset demand is significantly more affected by the information contained in the price and the large majority of trades are as predicted by empirical best response. Averaging over all situations where the naive benchmark differs from the empirical best response, the frequency of naive trading decisions is twice as high in SIM relative to SEQ, at 38% versus 19%.

Section 5 identifies various possible sources underlying the difficulty of hypothetical thinking in our markets. One possibility is that the participants feel rather well-informed by their own signals, relative to what they can learn from the price. We thus repeat the experiment with two treatments where early traders are much better informed than later traders, rendering learning from the price more important and more salient. We find that the replication only exacerbates the differences between simultaneous and sequential markets, both in terms of behavior and payoff consequences. This evidence makes it implausible that the bias is driven by negligence or the lack of salience of the price’s informativeness.

A further hypothesis is that the effect arises due to the difficulty in correctly interpreting human choices. As in the literature examining inference in games versus in single-person tasks (Charness and Levin, 2009; Ivanov et al., 2010), we therefore ask whether the bias also occurs if the price’s informativeness is generated by an automated mechanism. The corresponding treatment comparison replicates the main results. We can therefore rule out that the effect is driven by the necessity of responding to the behavior of others.

Finally, we ask whether the difficulty in contingent reasoning lies in the cognitive load of required inference, or rather in the hypothetical nature of price. To this end, we run another treatment where only one of the possible prices is considered, but still not yet realized. The rate of optimal choices in this treatment lies mid-way between that of the two main treatments, illustrating that the difficulty on contingent thinking is significantly fueled by both the amount and the hypothetical nature of possible prices in simultaneous markets.

We then combine the different treatments into an aggregate estimation of information use (Section 5.4). The analysis of the combined simultaneous treatments shows that relative to empirical best response, the participants under-weight the information contained in the price to a degree that is statistically significant (at $p = 0.09$ in a one-sided test) and that they strongly over-weight their own signal’s importance. In the sequential treatments, they over-weight both price and their own signal. Overall, the estimates indicate that traders far under-weight the prior distribution of the asset’s value but that they nevertheless learn too little from the price in simultaneous markets.

Taken together, the experiments provide evidence of an interaction between market microstructure and the efficiency of information usage. In the language introduced by Eyster and Rabin (2005), we find that the degree of ‘cursedness of beliefs’ is higher when the information contained in the price is less accessible: with price not yet realized, traders behave as if they tend to ignore the connection between other traders’ information and the price. Aggregate de-

mand therefore decreases too fast with the price. The economic bearing of the effect is further discussed in Section 6. We examine the predictions of Hong and Stein (1999) and Eyster et al. (2015) that markets with naive traders, who cannot learn from the price, generate an inefficient and slow price discovery. Naive traders tend to speculate against the price, pushing it back towards its ex-ante expectation also in cases where their own signals are consistent with the direction of price movement. Their erroneous speculation reduces the extent to which the price reveals the underlying value. Confirming this prediction, we simulate a standard price setting rule with our data and find that price discovery is slower in simultaneous treatments than in sequential treatments. Any (hypothetical) subsequent traders can therefore learn less from the price. But naiveté is detrimental not only to later players: also the observed payoffs of our market participants themselves are lower in SIM than in SEQ, albeit not to a large extent.

While we focus on markets, we again emphasize that our findings are also consistent with evidence in very different domains. The experimental literatures in economics and psychology provide several sets of related evidence that conditional inference is suboptimal. Psychologists have confirmed quite generally that decision processes depend on task complexity (Olshavsky, 1979) and that decision makers prefer decision processes with less cognitive strain. They focus on one model, one alternative or one relevant category when reflecting about possible outcomes and their consequences (Evans, 2007; Murphy and Ross, 1994; Ross and Murphy, 1996). They also process salient and concrete information more easily than abstract information (see e.g. Odean 1998 and the literature discussed there).

Several authors before us have pointed out that a possibility to reduce the complexity of learning is to proceed in a sequential mechanism, like in quote-driven markets.⁵ Our experiment suggests a specific manifestation of this effect, namely that drawing the attention to the realized price may enable the decision maker to interpret more easily the information underlying the price. In the related bilateral bargaining experiment by Carrillo and Palfrey (2011), buyers also trade more rationally in a sequential trading mechanism than in a simultaneous one. They process information more easily and exhibit less non-Nash behavior when facing a take-it-or-leave-it price instead of bidding in a double auction. Auction experiments similarly find that overbidding is substantially reduced in dynamic English auctions compared to sealed-bid auctions (Levin et al., 1996). Other contributions suggest that traders may systematically disregard relevant information that is conveyed by future, not yet realized events: overbidding decreases when finding the optimal solution does not necessitate updating on future events (Charness and Levin, 2009; Koch and Penczynski, 2014).⁶ Another related study is the voting experiment of Esponda and Vespa (2014) who find that when the voting rules follow a simultaneous game that requires hypothetical thinking, the majority of participants behave nonstrategically, whereas in the sequential design they are able to extract the relevant information from

⁵Shafir and Tversky (1992) note that participants see their preferences more clearly if they focus on one specific outcome. As they observe, "[t]he presence of uncertainty [...] makes it difficult to focus sharply on any single branch [of a decision tree]; broadening the focus of attention results in a loss of acuity" (p.457).

⁶Charness and Levin (2009) analyze the Winner's Curse in a takeover game, whereas Koch and Penczynski (2014) focus on common-value auctions.

others' actions and behave strategically.

We complement the described evidence on contingent thinking in strategic situations (bilateral bargaining games, auctions and strategic voting games) by addressing markets that clear exogenously and where traders are price takers. The simple structure of the traders' decision problems may make it easy for our participants to engage in contingent thinking—a possibility that the data refute—and helps us to straightforwardly assess whether market traders make too much or too little inference from the price.

2 Experimental Design

The basic framework is identical across treatments, involving a single risky asset and money. A market consists of two traders, trader 1 and trader 2, who each either buy or sell one unit of the risky asset.⁷ The asset is worth $\theta \in \{\underline{\theta}, \bar{\theta}\}$, with equal probabilities. Traders do not observe the fundamental value θ but they each receive a private signal $s_i \in [0, 1]$. The true value θ determines which of two triangular densities the signal is drawn from, such that in the low-value state the participants receive low signals with a higher probability, and vice versa:

$$f(s_i|\theta) = \begin{cases} 2(1-s_i) & \text{if } \theta = \underline{\theta} \\ 2s_i & \text{if } \theta = \bar{\theta} \end{cases} \quad i \in \{1, 2\} \quad (1)$$

Conditional on θ , the signals of the two traders are independent.

Each trader i faces a separate transaction price p_i . Trader 1's price p_1 is uniformly distributed in $[\underline{\theta}, \bar{\theta}]$ and is uninformative about the fundamental value θ . Trader 1 observes his private signal s_1 and states his maximum willingness to pay by placing a limit order b_1 . If p_1 lies weakly below b_1 , he buys one unit of the asset. If p_1 strictly exceeds b_1 , he sells one unit.⁸ By checking an additional box, trader 1 may convert his limit order into a “reversed” limit order. A reversed limit order entails the opposite actions: the trader buys if p_1 weakly exceeds b_1 , otherwise he sells. (Only few participants make use of it; we defer the motivation for allowing reversed limit orders to Section 2.2.) Let Z_1 denote the indicator function that takes on value 1 if a limit order is reversed. Trader 1's demand is $X_1(p_1, b_1)$:

$$X_1(p_1, b_1) = Y_1(p_1, b_1)(1 - Z_1) - Y_1(p_1, b_1)Z_1 \quad (2)$$

$$Y_1(p_1, b_1) = \begin{cases} 1 & \text{if } p_1 \leq b_1 \\ -1 & \text{if } p_1 > b_1 \end{cases} \quad \text{where } p_1 \sim U[\underline{\theta}, \bar{\theta}]$$

The task of trader 2 varies across the two main treatments, a simultaneous and a sequential mechanism.

⁷Because of a possible reluctance to sell short, we avoid any notion of short sales in the experimental instructions. Participants are told that they already possess a portfolio that needs to be adjusted by selling or buying one unit of a given asset.

⁸The design does not allow for a “no trade” option because of the possibility that it may add noise and complications to the data analysis. We opted for a minimal set of actions that enables participants to state their preference to buy and sell with a single number.

2.1 Simultaneous treatment (SIM)

Trader 2 observes trader 1's price p_1 and her own private signal s_2 . Like trader 1, she chooses a limit order or, optionally, a reversed limit order. When submitting her decision, she does not know her own price p_2 .

Participants are informed that the price p_2 reflects the expectation of an external market maker who observes trader 1's buying or selling decision and who assumes that trader 1 bids rationally upon receipt of his signal s_1 . Importantly, to avoid any ambiguity in the description, they learn the pricing rule that maps p_1 and the realized value of X_1 into p_2 :

$$p_2 = \begin{cases} \frac{\bar{\theta} + p_1}{2}, & \text{if } X_1(p_1, b_1) = 1 \\ \frac{\underline{\theta} + p_1}{2}, & \text{if } X_1(p_1, b_1) = -1 \end{cases} \quad (3)$$

Participants also receive a verbal explanation of the implied fact that for given p_1 , trader 2's price p_2 can take on only one of the two listed possible realizations, depending on whether trader 1 buys or sells. Through X_1 , p_2 is influenced by trader 1's private signal s_1 ; p_2 is therefore informative about the asset value θ and trader 2 would ideally condition her investment decision on both s_2 and p_2 .

2.2 Sequential treatment (SEQ)

In treatment SEQ, trader 2 observes the price p_2 as specified in (3) *before* making her decision. The game proceeds sequentially, with trader 1 first choosing his (possibly reversed) limit order b_1 . As in treatment SIM, his demand X_1 determines the price for trader 2, p_2 . Trader 2 observes the realized value of $\{p_1, p_2, s_2\}$ and chooses between buying and selling at p_2 .

It is straightforward to check that treatments SIM and SEQ are strategically equivalent. Treatment SEQ allows for four possible strategies contingent on $p_2 \in \{\frac{\underline{\theta} + p_1}{2}, \frac{\bar{\theta} + p_1}{2}\}$: $\{buy, buy\}$, $\{buy, sell\}$, $\{sell, buy\}$ and $\{sell, sell\}$. In treatment SIM, the possibility to reverse the limit order enables the same four combinations of buying and selling contingent on p_2 . The two strategy spaces are therefore isomorphic.

2.3 Payoffs

In each of the treatments, the experimenter takes the other side of the market, which therefore always clears. In case of a buy, the profit Π_i of trader $i \in \{1, 2\}$ is the difference between the asset value and the market price, and vice versa if the asset is sold:

$$\Pi_i = (\theta - p_i)X_i(p_i, b_i) \quad (4)$$

Between treatments SIM and SEQ, payoffs arising from each combination of strategies and signals are identical. Any rational response to a fixed belief about trader 1 leads to the same purchases and sales in the two treatments.⁹

⁹This statement holds under the assumptions of subjective utility theory. Probability weighting and other generalizations of expected utility can lead to different weighting of outcomes between the two treatments.

3 Predictions

We mainly focus on trader 2 and compare the participants' behavior to three theoretical predictions. The first two are variants of the case that trader 2 has rational expectations and properly updates on her complete information set. As the third benchmark, we consider the case that trader 2 fully neglects the price's informativeness. For all three predictions, we assume traders to be risk-neutral.

3.1 Rational best response

Trader 1 has only his private signal s_1 to condition his bid upon. His optimal limit order b_1^* is not reversed and maximizes the expected profit conditional on s_1 . It is easy to show (using the demand function (2)) that b_1^* increases linearly in the signal:

$$b_1^*(s_1) = \arg \max_{b_1} E[(\theta - p_1)X_1(p_1, b_1)|s_1] = E[\theta|s_1] = \underline{\theta} + (\bar{\theta} - \underline{\theta})s_1 \quad (5)$$

Under rational expectations about trader 1's strategy, trader 2 maximizes her expected payoff conditioning on both her private signal s_2 and the informative price p_2 . If her maximization problem has an interior solution, it is solved by the following fixed point:¹⁰

$$b_2^*(s_2) = E[\theta|s_2, p_2 = b_2^*(s_2)] \quad (6)$$

The optimal bidding of trader 2 never uses reversed limit orders but follows a cutoff strategy that switches from buying to selling as the price increases. At a price equal to the (interior) cutoff b_2^* , the trader is indifferent between a buy and a sell.

The Bayes-Nash (BN) strategy of trader 2, however, simplifies to a step function: p_2 reflects the market maker's expectation (see (3)), implying that in equilibrium p_2 would make trader 2 indifferent in the absence of her own signal s_2 . The additional information contained in s_2 breaks the tie, such that trader 2 buys for $s_2 \geq E[s_2] = \frac{1}{2}$, and sells otherwise.

However, the BN best response is not the most payoff-relevant 'rational' benchmark. In the experiment, participants in the role of trader 1 deviate from their best response b_1^* and participants acting as trader 2 would optimally adjust to it. Their price p_2 is still informative about θ because it reflects s_1 , but p_2 does not generally equal $E[\theta|X_1^*(\cdot)]$ if X_1 is subject to deviations from $X_1^*(p_1, b_1^*)$. For a stronger test of naive beliefs, we therefore consider the empirical best response (EBR) to the participants acting as traders 1. The empirical best response is computed via a numerical approximation to the fixed point equation (6).

The two benchmarks BN and EBR are depicted in Figure 1 (for the parameters of the actual experiment that are reported in Section 4, and using the empirical behavior described in Section 5 for the calculation of EBR), together with the naive prediction that we describe next.¹¹ The graphs represent the

¹⁰For a simple proof of this statement, verify that if b_2^* were to violate (6) then there would exist realizations of (p_2, s_2) such that p_2 lies in the vicinity of $E[\theta|s_2, p_2 = b_2^*]$ and profits are forgone.

¹¹The kinks in the EBR function arise because of the numerical approximation to the fixed point, which is done for signals that are rounded to lie on a grid with step size 0.1 for close approximation.

prices at which, for a given signal, trader 2 is indifferent between buying and selling. She is willing to buy at prices below the graph and willing to sell at prices above the graph. The EBR graph is less steep than that of BN: e.g. for an above-average level of p_2 , EBR requires trader 2 to buy only if she has additional positive information (large s_2).

3.2 Best response to naive beliefs

Contrasting the optimal behavior, a trader 2 with naive beliefs does not infer any information from the price. She fails to account for the connection between trader 1's signal s_1 and his demand X_1 and, instead, conditions on her own signal s_2 only. The maximization problem with naive beliefs is then analogous to that of trader 1 and leads to the same bidding behavior:

$$b_2^N = \arg \max_{b_2} E[(\theta - p_2)X_2(p_2, b_2)|s_2] = E[\theta|s_2] = \underline{\theta} + (\bar{\theta} - \underline{\theta})s_2 \quad (7)$$

The naive strategy is depicted as the straight line in Figure 1. Its underlying belief is equivalent to level-1 reasoning or fully cursed beliefs. In the level-k framework (for a formulation with private information, see e.g. Crawford and Iriberri (2007)) level-0 players ignore their information and randomize uniformly and a naive trader 2, as defined above, is therefore equivalent to a level-1 agent. In our setting, this prediction also coincides with a fully cursed strategy of Eyster and Rabin (2005) and Eyster et al. (2015) that best responds to the belief that agent 1's equilibrium mixture over bids arises regardless of their information.¹²

3.3 Hypotheses

As outlined in the Introduction, we conjecture that the updating on additional market information is more difficult in the simultaneous treatment than in the sequential treatment. Using the benchmarks from the previous subsection, we translate the conjecture into a behavioral hypothesis:

Hypothesis 1 *Naive bidding is more prevalent in treatment SIM than in treatment SEQ.*

The hypothesis is tested in the next section by considering those decisions of trader 2 where EBR and Naive bidding differ, separately for each of the two treatments. As shown in Figure 1, EBR and Naive bidding predict different decisions in the area between the two graphs. For instance, at prices within this area, a naive agent with a signal below 0.5 would buy whereas she would sell according to EBR.

¹²In fully cursed equilibrium, trader 2 believes that trader 1 with signal s_1 randomizes uniformly over his possible bids: trader 2 expects that trader 1 with signal s_1 has a bid distribution equal to that resulting from the optimal bids given in (5), independent of s_1 . The perceived mixture of bids by each type of trader 1 therefore follows the distribution $F(\frac{b-\underline{\theta}}{\bar{\theta}-\underline{\theta}}) = F(s_1)$, with density $\frac{1}{2}f(s_1|\underline{\theta}) + \frac{1}{2}f(s_1|\bar{\theta}) = 1$. The analysis of Eyster and Rabin (2005) and Eyster et al (2015) also allows for intermediary levels of cursedness, where agents may only partially ignore the information revealed by other agents' actions. Our estimations in Subsection 5.4 also allow for milder versions of information neglect.

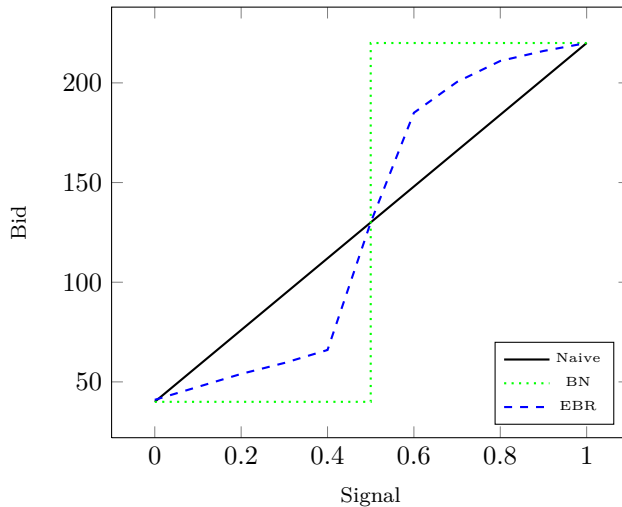


Figure 1: Naive, Bayes Nash and empirical best responses of trader 2.

Our second hypothesis considers the possibility that *all* participants acting as trader 2 have naive beliefs. In this case, the symmetry of the two traders' decision problems would induce symmetry between their bid distributions. We can therefore use trader 1's bid distribution as an empirical benchmark for naive traders 2. We restrict the comparison to treatment SIM, where the two traders have identical action sets.

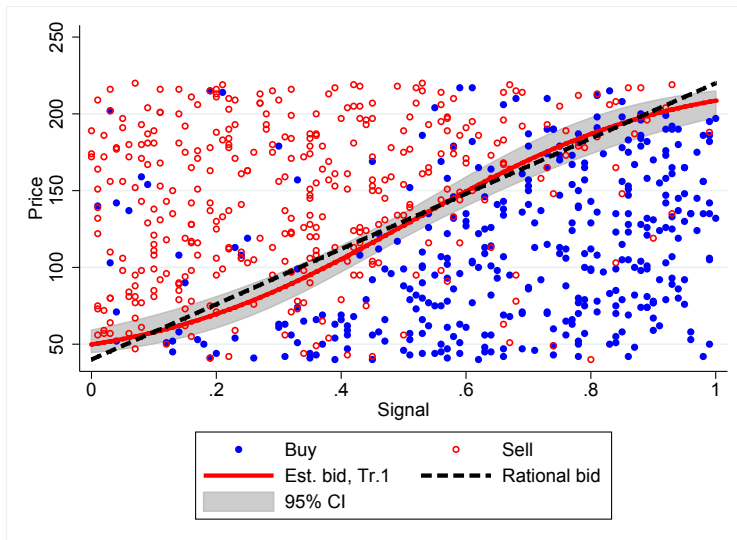
Hypothesis 2 *In treatment SIM, bids of trader 2 do not significantly differ from bids of trader 1.*

4 Experimental Procedures and Results

4.1 Procedures

The computerized experiment is conducted at Technical University Berlin, using the software z-Tree (Fischbacher, 2007). A total of 144 students are recruited with the laboratory's ORSEE database (Greiner, 2004). 72 participants are in each of the treatments SIM and SEQ, each with three sessions of 24 participants. Within each session, the participants are divided into two equally sized groups of traders 1 and traders 2. Participants remain in the same role throughout the session and repeat the market interaction for 20 periods. At the beginning of each period, participants of both player roles are randomly matched into pairs and the interaction commences with Nature's draw of θ , followed by the market rules as described in Section 2. At the end of each period, subjects learn the value θ , their own transaction price (if not already known) and their own profit. Upon conclusion of the 20 periods, a uniform random draw determines for every participant one of the 20 periods to be paid out for real.

Participants read the instructions for both roles, traders 1 and 2, before learning which role they are assigned to. The instructions include an elaborate computer-based simulation of the signal structure as well as an understanding test. The



Note: The average bidding curve corresponds to $\theta + (\bar{\theta} - \theta) \cdot \hat{P}(X_1|s_1)$, where $\hat{P}(X_1|s_1)$ is the probit estimate of the probability of trader 1 buying in treatment SIM.

Figure 2: Trades of traders 1.

support of the asset value is $\{40, 220\}$.¹³ Each session lasted approximately 90 minutes and participants earned on average EUR 22.02. Total earnings consist of a show-up fee of EUR 5.00, an endowment of EUR 15.00 and profits from the randomly drawn period (which could be negative but could not deplete the entire endowment). Units of experimental currency are converted to money by a factor of EUR 0.08 per unit.

4.2 Results

4.2.1 Trader 1

For a cleaner comparison of the two treatments, we analyze realized trades instead of bids, thereby considering also the suboptimal, reverse limit orders (ca. 15% of all bids in treatment SIM). Figure 2 shows the implemented buys and sells of participants acting as trader 1 in treatment SIM, with the corresponding market price on the vertical axis and their private signal on the horizontal axis. (Results for trader 1 in treatment SEQ are very similar.) The figure also includes the theoretical prediction and the results of a probit estimate of the mean bid. The mean bid increases in the signal, even slightly stronger than is predicted by the benchmark theory. This overreaction is not significant, though.

4.2.2 Trader 2: Testing hypotheses

Hypothesis 1. To evaluate the degree of naiveté, we focus on the area of Figure 1 where naive and optimal strategies make different predictions. That is, we

¹³See the Online Appendix for a set of instructions for treatments SIM and SEQ. We chose the possible asset values $\{\theta, \bar{\theta}\} = \{40, 220\}$ in an attempt to minimize midpoint effects, which are often observed in experiments.

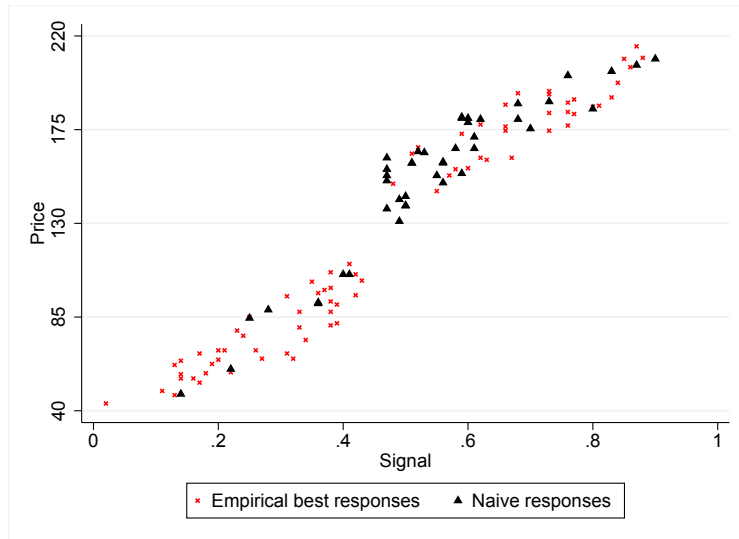


Figure 3: Sells and buys within the relevant area in treatment SIM.

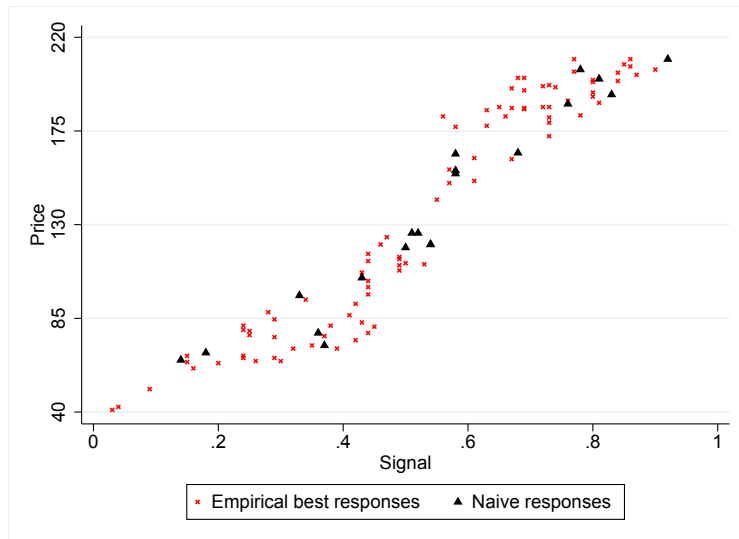


Figure 4: Sells and buys within the relevant area in treatment SEQ.

consider the set of trades with prices and signals between the solid (b_2^{Naive}) and the dashed (b_2^{EBR}) bidding functions: $\{(s_2, p_2) | (s_2 < 1/2) \cap (b_2^{EBR} \leq p_2 \leq b_2^{Naive})\} \cup \{(s_2, p_2) | (s_2 > 1/2) \cap (b_2^{Naive} \leq p_2 \leq b_2^{EBR})\}$. Within this area, we calculate the proportion η of naive decisions:

$$\eta = \frac{d_N}{d_N + d_B} \quad (8)$$

where d_N and d_B denote the number of orders consistent with naive and EBR predictions, respectively.

Figures 3 and 4 show the relevant observations in treatments SIM and SEQ, respectively. For these observations, naive expectations induce buys for signals below 0.5 and sells for signals above 0.5, while rational expectations induce opposite actions. The empirical measures d_N and d_B correspond to the number of triangle markers and cross markers, respectively. Hypothesis 1 is confirmed if the proportion of naive choices is larger in treatment SIM than in treatment SEQ: $\eta^{SIM} > \eta^{SEQ}$.

Indeed, we find that neglect of information contained in the price is stronger in a simultaneous market. Appendix Table A4 shows that the share of naive decisions in treatment SIM ($\eta = 0.38$) is twice as large as in treatment SEQ ($\eta = 0.19$). The difference is statistically significant ($p = 0.0091$, Wald test).

An especially strong difference between the two treatments appears in situations where trader 2 has a relatively uninformative signal, $s_2 \in [0.4, 0.6]$, i.e. when traders have the strongest incentive to make trading contingent on the price. In these cases, the frequency of buying at a price below the ex-ante mean of $p_2 = 130$ is at 0.68 in SIM and at 0.37 in SEQ (see Appendix Table A1). Similarly, the frequency of buying at a high price, above $p_2 = 130$, is at 0.28 in SIM and at 0.48 in SEQ (see Appendix Table A2). This illustrates that treatment SEQ's participants were less encouraged by low prices and less deterred by high prices, respectively, than treatment SIM's participants, consistent with a relatively more rational inference in the sequential market.

In Appendix A.3, we also consider the evolution of decisions in the course of the experiment. We cannot detect any learning success over 20 repetitions.¹⁴

Hypothesis 2. Hypothesis 2 compares the buy and sell decisions of traders 1 and 2 in treatment SIM. Figure 5 reveals that the two traders' average bid functions do not significantly, or even perceptibly, differ from each other. Just like trader 1, trader 2 shows no significant deviations from a linear bidding function, an observation that is consistent with full naiveté of trader 2.¹⁵

We note that in the variations of the simultaneous game, featuring in the next section, fully naive bidding does not always appear.

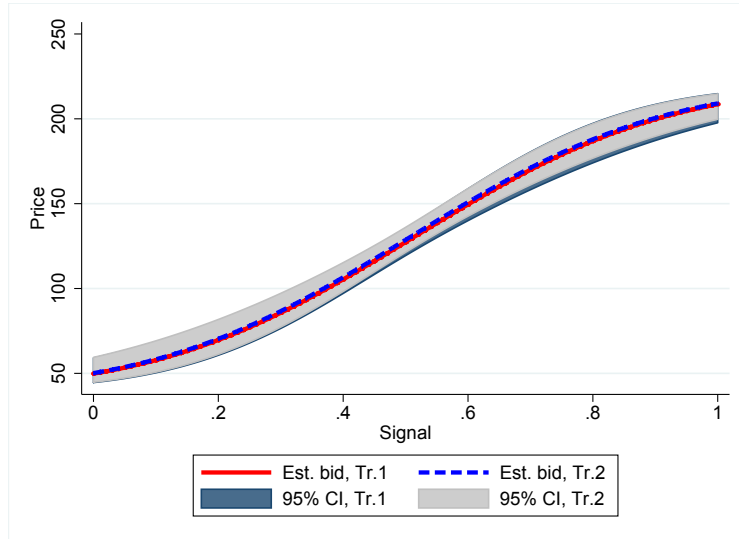


Figure 5: Estimated average bids of traders 1 and 2 in treatment SIM.

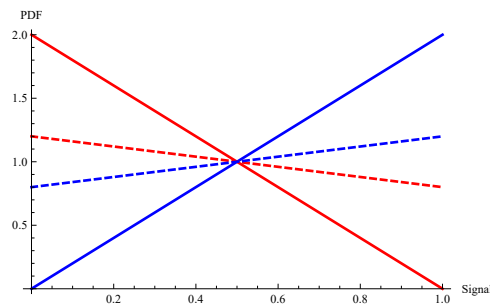


Figure 6: Signal distributions for trader 1 (solid) and trader 2 (dashed) in LSQ treatments.

5 Possible drivers of information neglect

5.1 Signal strength

One possible driver of the observed information neglect is that the participants' strong private signals might distract them from the information contained in the price. In a challenging and new environment, participants may perceive the benefit from interpreting the price as relatively low. In real markets, investors may be more attentive to the price's informativeness, especially when they themselves have little private information.¹⁶

We examine the hypothesis by introducing an asymmetric signal strength between trader 1 and trader 2, keeping the rest of the design unchanged. In two additional treatments with "Low Signal Quality", LSQ-SIM and LSQ-SEQ (with $N = 70$ and $N = 68$, respectively), trader 2's signal is less informative. The densities in the new treatments are depicted in Figure 6 and take the following form.

$$\begin{aligned}f(s_i|\theta = \underline{\theta}) &= 1 - \tau_i(2s_i - 1) \\f(s_i|\theta = \bar{\theta}) &= 1 + \tau_i(2s_i - 1) \\&\text{with } \tau_1 = 1 \text{ and } \tau_2 = 0.2.\end{aligned}$$

Behavior of trader 2 deviates from the naive prediction in both treatments LSQ-SIM and LSQ-SEQ. Trader 2s react to their signals more strongly than predicted by naive bidding (see Figure A 1). A comparison with the bids in the main treatments SIM and SEQ thus supports the conjecture that subjects pay more attention to market information when they are less informed privately.

However, the discrepancy between the two market mechanisms increases with information asymmetry. The share of naive decisions in treatment LSQ-SEQ (22%, black triangles in Figure 7b) is much smaller than in LSQ-SIM (44%, black triangles in Figure 7a). This significant difference ($p = 0.0003$, Wald test) corresponds to a steeper estimate of the average bidding curve in LSQ-SEQ, see Appendix Figure A 1. Tables A1 to A3 in the appendix also show that differences in frequencies of buys and sells between the two mechanisms are highly significant for various signal ranges, and that they tend to be larger than in the comparison of SIM and SEQ. For example, participants in the role of trader 2 of LSQ-SEQ act very frequently against their own signal. In sum, the importance of trading mechanisms for rational decision making prevails under the new informational conditions.

5.2 Strategic uncertainty

Strategic uncertainty adds to the complexity of the trading game. For an accurate interpretation of price, participants in the role of trader 2 need to consider the trading behavior of trader 1 and their ability to do so may vary between

¹⁴Carrillo and Palfrey (2011) report similar evidence of constantly naive play in their experiment.

¹⁵In contrast, there do appear significant differences from naive actions in treatment SEQ, which is in line with the previously examined Hypothesis 1. Results are available upon request.

¹⁶We thank an anonymous referee for raising this hypothesis.

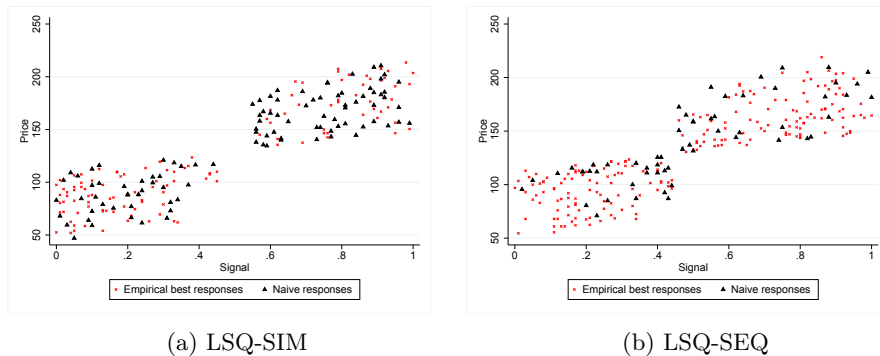


Figure 7: Buys and sells consistent with either naive bidding or EBR in treatments LSQ-SIM and LSQ-SEQ, respectively.

simultaneous and sequential mechanisms. In other words, the necessity to assess the human-driven EBR (not just the simpler BN response) may lead to less optimal behavior by trader 2 in treatment SIM relative to SEQ.

We therefore examine whether the treatment effect appears also in two additional treatments with “No Player 1” (NP1), containing 40 participants in NP1-SIM and 46 in NP1-SEQ, all of whom act in the role of trader 2. In these treatments we delete trader 1’s presence. Participants acting as trader 2 are informed that the price is set by a market maker who receives an additional signal. This additional signal follows a distribution that mimics the information of the market maker in the two main treatments when observing the demand X_1 of a trader 1 who behaves rationally.¹⁷

For better comparison with the main treatments, the instructions of the NP1 treatments retain not only much of the wording but also the chronological structure of the main treatments. Participants in NP1 treatments thus learn about the existence of p_1 , which is presented to them as a random “initial value” of the asset’s price, and they learn that the market maker observes an additional signal that is correlated with the asset’s value. Like in the main treatments, the instructions display the updating rule (3) and explain that it results in the price p_2 at which the participants can trade and which reflects the expectation of the asset’s value, conditional on the market maker’s additional signal but not conditional on the participants’ own signal.

The data show no strong differences between the NP1 treatments and the main treatments. Appendix Figure A 2 shows that the estimated bidding curve in NP1-SIM exhibits the same slope as the curve in SIM, with a mild downward shift, whereas behavior in NP1-SEQ is very close to that of SEQ.¹⁸

Most notably, the effect of simultaneous versus sequential trading persist. The share of naive decisions is two and a half times higher in NP1-SIM than

¹⁷The distributions of the additional signals (one for each asset value) are shown in a graphical display. The instructions do not explain how the distributions are determined.

¹⁸The downward shift in NP1-SIM is more pronounced for low signals and leads to a significant deviation from the naive benchmark (Multiple binomial testing with Bonferroni correction rejects 1 out of 9 hypotheses at .0055 significance level, see Appendix A.2). Despite this deviation, the average bid does not increase disproportionately in the private signal as the rational benchmark predicts. Another mild difference is that the use of reversed limit orders is smaller in NP1-SIM (9%) than in SIM (15%).

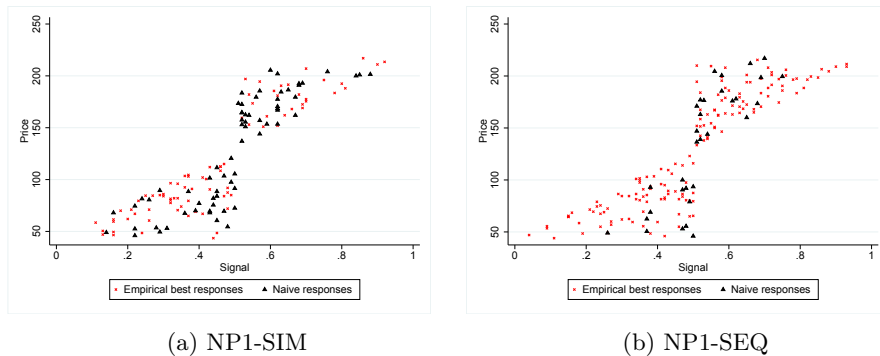


Figure 8: Naive v Bayesian in NP1-SIM and NP1-SEQ.

in NP1-SEQ (45.27% vs. 17.67%). We also observe significantly more buys at high prices and more sells at low prices in NP1-SEQ (see Tables A1 and A2 in Appendix A.1). Figure 8 shows the individual decisions for cases where naive and rational predictions differ, in treatments NP1-SIM and NP1-SEQ, respectively.

5.3 Number of decisions per treatment

Our last treatment addresses the question whether the higher frequency of naive decisions in SIM may be driven by the additional cognitive strain that conditional thinking requires. Perhaps, it is not conditionality per se that is difficult for the participants, but rather the fact that they have to make two decisions in treatment SIM (one for each possible price realization) but only one in treatment SEQ.

We therefore introduce a “hypothetical” sequential treatment (Hyp-SEQ) with 62 participants, which rules out higher dimensionality of strategies as a source of difficulty. Treatment Hyp-SEQ is analogous to SEQ in that after learning trader 1’s price p_1 , participants in the role of trader 2 specify their buying or selling preferences for only a single price \hat{p}_2 . However, \hat{p}_2 is only a candidate price as \hat{p}_2 is equiprobably drawn from the two price values that are possible after updating via rule (3). Participants decide whether they would buy or sell at \hat{p}_2 and the decision is implemented if and only if trader 1’s demand induces the realization $p_2 = \hat{p}_2$. Otherwise, trader 2 does not trade and makes zero profit.

Participants in treatment Hyp-SEQ thus face only one price and make only one decision, rendering the task dimensionality identical to that in SEQ. (The instructions are almost word-for-word identical.) But the nature of the decision in Hyp-SEQ is conditional, like in treatment SIM. We can therefore assess the importance of task dimensionality by comparing SIM versus Hyp-SEQ, and the role of conditionality by comparing SEQ versus Hyp-SEQ.

Average bidding shows no large difference between treatments SIM and Hyp-SEQ, or between traders 1 and 2 of treatment Hyp-SEQ: The estimated bid functions in Appendix Figures A 3a and A 3b exhibit approximately the same slope. Moreover, the Appendix A.2 also shows that naive bidding cannot be rejected for treatment Hyp-SEQ, in multiple binomial testing.

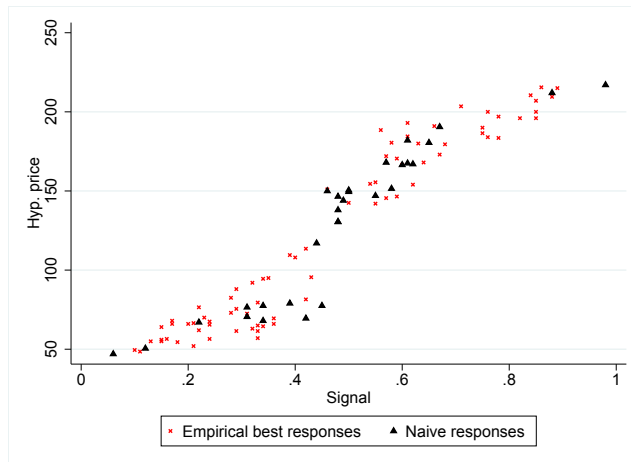


Figure 9: Naive v EBR in Hyp-SEQ.

However, Figure 9 and Table A4 in the Appendix show that the frequency of making suboptimal decisions (η) in Hyp-SEQ lies well in between those of SEQ and SIM. The significant difference between treatments SIM and Hyp-SEQ (0.38 versus 0.28, $p=0.022$, one-sided t test) shows that reducing the set of hypothetical prices considerably improves decision-making. Yet, the frequency of naive decisions is still significantly higher in Hyp-SEQ than in the fully sequential treatment SEQ, (0.28 versus 0.19, $p=0.081$, one-sided t test).¹⁹ Altogether, we conclude from the above tests that reducing the number of hypothetical trading decisions reduces the degree of naiveté, but does not eliminate it.²⁰

5.4 Random Utility Model

This subsection pools the data for a statistical comparison of sequential versus simultaneous mechanisms. We combine the data from all simultaneous treatments into a data set “SIM+” and those from sequential treatments into a data set “SEQ+”. (Data from the hybrid treatment Hyp-SEQ are excluded.) We assume that the probability with which trader 2 buys the risky asset follows a logistic distribution, allowing for an over-weighted or under-weighted relevance of the available pieces of information:

$$P(X_2 = 1 | u_i, s_2, p_2) = \frac{e^{\lambda(\hat{E}[\theta | p_2, s_2] - p_2 + u_i)}}{1 + e^{\lambda(\hat{E}[\theta | p_2, s_2] - p_2 + u_i)}} \quad (9)$$

¹⁹Notice that the lower rate of suboptimal decisions in Hyp-SEQ relative to SIM is consistent with the main idea of Li’s (2016) obvious strategy proofness: in Hyp-SEQ, the set of relevant prices is reduced to a singleton, helping the participants to detect the optimal strategy.

²⁰Our working paper version, Ngangoue and Weizsäcker (2015) shows a first version of the experiment where the simultaneous treatment elicits buy and sell preferences for a list of 26 hypothetical prices (treatment “Price List”), instead of 2 as in the present paper’s treatment SIM. There, we find the neglect of the price informativeness to be even more pronounced, which is also consistent with an effect of task dimensionality. The previous experiment, however, also has other differences to the present one.

with

$$\hat{E}[\theta|p_2, s_2] = 40 + 180 \cdot \hat{P}(\theta = 220|p_2, s_2) \quad (10)$$

$$\hat{P}(\theta = 220|p_2, s_2) = [1 + LR(s_2)^{-\beta} \cdot LR(p_2)^{-\alpha}]^{-1} \quad (11)$$

The choice probability (9) depends on subjectively expected payoff, $\hat{E}[\theta|p_2, s_2] - p_2$. The parameter λ reflects the precision of the logistic response and u_i is the random utility shifter, which we assume to be normally distributed with mean 0 and variance σ_u^2 . To allow for irrational weighting of information, we introduce the subjective posterior probability of the event that $\theta = 220$, given by $\hat{P}(\theta = 220|p_2, s_2)$. Analogous to the method introduced by Grether (1992), we let the posterior probability depend on the likelihood ratios of the signal and the price, $LR(s_2) \equiv \frac{P(\theta=220|s_2)}{P(\theta=40|s_2)}$ and $LR(p_2) \equiv \frac{P(\theta=220|p_2)}{P(\theta=40|p_2)}$, respectively. The likelihood ratios are exponentiated by the potentially irrational weights β and α that the participant assigns to the signal's and the price's informational content. A participant with naive beliefs (a 'fully cursed' participant) would correctly weight the signal, $\beta = 1$, but would ignore the information in the price, $\alpha = 0$. An intermediary level of cursedness translates into α between 0 and 1. A rational trader would correctly weight the signal and the price, $\beta = \alpha = 1$. The model also allows for an over-weighting of the signal or the price, by letting β or α exceed 1.

We estimate the model via Maximum Simulated Likelihood (MSL). To arrive at $LR(p_2)$, we estimate the distributions $P(p_2|\theta = 220)$ and $P(p_2|\theta = 40)$ for each treatment individually via kernel density estimation and infer $\frac{P(\theta=220|p_2)}{P(\theta=40|p_2)}$ for each p_2 in the data set.

Table 1: Results of MSL estimation

	Trader 1	Trader 2	
		SIM+	SEQ+
β	2.05*** (0.31)	2.54** (0.90)	1.36*** (0.36)
α	-	0.60* (0.26)	1.85*** (0.22)
λ	0.0314*** (0.003)	0.0230*** (0.004)	0.0373*** (0.006)
σ_u	0.0010	0.0010	0.0039
N	3435	2220	2260

Note: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Std. Err. in parentheses. Hypothesis testing for β and α refers to one-sided tests of deviations from 1. The estimation for trader 1 pools all treatments with participants acting as trader 1 since their data do not significantly differ across treatments.

The estimates are reported in Table 1 and confirm the findings of the previous subsections. Trader 1's model estimates serve as a benchmark. Participants in

the role of trader 1 overweight their private signal ($\beta = 2.05$), inducing a slight S-shape of the estimated bid function (see Figure A 4). Traders' 2 weighting of the private signal decreases from 2.54 to 1.36 between the simultaneous and the sequential treatments. Both of these β estimates significantly differ from 1, but in the sequential treatments β lies significantly below trader 1's weighting of the private signal ($p = 0.0298$, Wald test).

In the simultaneous mechanisms, the estimated α of 0.60 lies well below the optimal value 1, albeit at a somewhat marginal statistical significance of $p = 0.09$. While this difference from 1 reflects the hypothesis that participants pay too little attention to the price's informativeness, we can also reject the extreme formulation of Hypothesis 2, stating that participants are fully naive: α differs significantly from 0.

In the treatments with sequential mechanisms, the perceived levels of informativeness of signal relative to price are reversed. These treatments induce a significant over-weighting of the price's likelihood ratio ($\alpha = 1.85$).²¹ Overall, the evidence from sequential treatments shows that the prior distribution of θ is under-weighted and that, confirming Hypothesis 1, sequential markets reveal a significantly stronger inference from the price than simultaneous markets.

6 Discussion: Information neglect in markets

This section discusses the possible impact of naiveté on market efficiency. We begin by stating a classical question of market prices: how do prices that arise after a given trading pattern differ from equilibrium prices? Notice that this question addresses the welfare of *subsequent* traders in the same market, i.e., traders outside of the set of traders that we consider in the experiment. We therefore have to resort to auxiliary calculations. Yet we also consider the payoff of our actual participants.

Pricing. A natural measure of price efficiency is the speed at which price aggregates the traders' dispersed pieces of information and converges to fundamental value. With naive traders in the market, this speed may be reduced. Moreover, naive traders may distort the price recovery process by suppressing some subsets of possible signals more than others. Two theoretical contributions that study the implications of naiveté on price are by Hong and Stein (1999) and Eyster et al. (2015). They both find, with different models, that the presence of naive traders creates a bias of prices leaning towards their ex-ante expectation. The reason is that naive traders are likely to engage in excessive speculation based on their own signal—they bet against the market price too often. This pushes price towards its ex-ante mean.²²

Testing this implication requires the simulation of a specific price mechanism after trader 2 has completed her trades. For simplicity and for consistency with the rule governing p_2 , we calculate the price that a market maker would

²¹This relates to Levin et al. (1996)'s finding that participants in the English auction put relatively more weight on the latest drop-out prices compared to their own signal.

²²Hong and Stein (1999) analyze a dynamic model where information dispersion is staggered in the market and where naive traders are myopic but can be exploited by sophisticated (yet cognitively restricted) traders who start betting against the naive traders eventually. Price can therefore overshoot at a later stage in the cycle. Eyster et al.'s (2015) model uses partially cursed equilibrium to show the bias in pricing, using a more standard (and more static) model of financial markets with incomplete information akin to that in Grossman (1976).

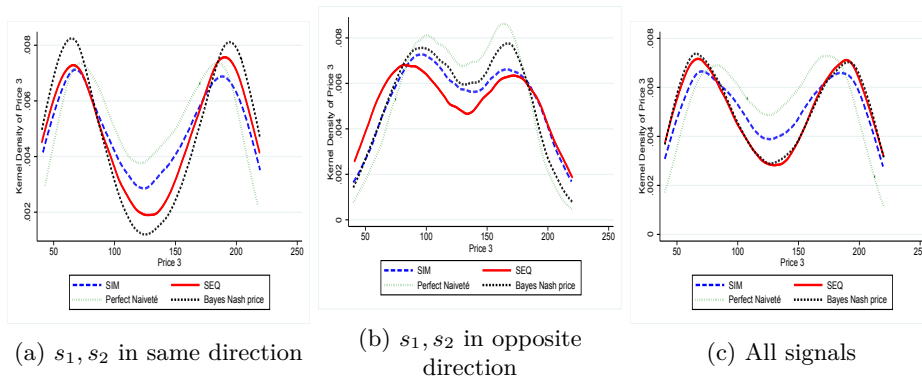


Figure 10: Kernel density of efficient price 3 after naive, rational and actual demand of traders 1 and 2 in SIM and SEQ .

set in Bayes Nash equilibrium: the market maker sets the price p_3 equal to $E[\theta|X_1, X_2]$, where $X_1, X_2 \in \{-1, 1\}$ denote the demand of traders 1 and 2, assumed to follow the Bayes-Nash prediction. In our main treatments SIM and SEQ, the price for a hypothetical trader 3 is thus a simple function of p_2 and X_2 :²³

$$p_3 = \begin{cases} \frac{-8800+310p_2}{50+p_2} & \text{if } X_2 = 1 \\ \frac{-8800+50p_2}{310-p_2} & \text{if } X_2 = -1 \end{cases}$$

Under the given pricing rule, price moves towards its extremes fast if both signals s_1 and s_2 deviate from their expectation in the same direction. In this case either both traders buy or both traders sell, in Bayes Nash Equilibrium. For all cases where s_1 and s_2 lie on the same side of 0.5, Figure 10a shows the resulting distribution of Bayes Nash price p_3 as a dotted line, with much probability mass located towards the extremes. In contrast, if trader 2 bids naively, then she will tend to sell at high prices and buy at low prices, creating excessive density of p_3 near the center of the distribution (light grey line).

Figure 10a also depicts the kernel densities of the price p_3 that would arise from the actual trading in treatments SIM and SEQ. The price distribution under SIM is close to that of naive bidding. In SEQ, prices deviate more from the prior expectation and the distribution lies far closer to its equilibrium prediction.

Figure 10b shows the kernel densities when the two signals are on opposite sides of their ex-ante expectation. Here, the aggregate information is not very informative, prices with naive and Bayes-Nash traders do not differ much and markets yield prices that revolve around prior expectations. Figure 10c depicts the densities when taking into account all observations. Overall, the price distribution in treatment SEQ has a more pronounced bi-modal shape.

In a nutshell, prices in the simultaneous mechanisms incorporate information slowly. This finding is consistent with the momentum effect in call auctions documented in Amihud et al. (1997) and Theissen (2000).

To quantitatively assess price efficiency under the two treatments, we ask about the variance of fundamental value conditional on the price, $Var[\theta|p_3]$.

²³In treatments LSQ, we obtain $p_3 = \frac{1030(-8.54p_2)}{770+p_2}$ if $X_2 = 1$, $p_3 = \frac{-770(11.43p_2)}{p_2-1030}$ else.

It captures the error in market expectations given information contained in p_3 . Conditional variance is significantly lower in treatment SEQ than in SIM, at high level of significance ($p=0.00$, nonparametric median test, taking each market as a unit of observation) and with a somewhat sizable difference: in treatment SIM, the price explains on average 21% of the variance in the asset value, versus 27% in treatment SEQ.²⁴

Profits. The difference between simultaneous and sequential mechanisms also affects the distribution of profits of trader 2. A corresponding difference occurs in each of the relevant treatment comparisons, but it is economically small (our experiments were not designed to generate big payoff differences between treatments) and is statistically significant only in the comparison LSQ-SIM versus LSQ-SEQ, i.e. with asymmetry in the informativeness of signals. Less informed traders benefit from sequential information processing, where the employed updating is more rational. The results on mean and median profits of each treatment is in Table A5 in the Appendix. It is also noteworthy that the distribution of profits conditional on price p_2 in LSQ-SEQ is mirror-inverted to the one in LSQ-SIM (see Figure A 5b): the majority of traders in LSQ-SIM lose significant amounts, whereas the majority of traders in LSQ-SEQ make gains. This hints at the importance of pre-trade transparency to restrain insider trading in real-world markets. Naive later traders may suffer if they are poorly informed.

Trading volume. Naive beliefs may not only affect prices and profits, but may also trigger speculative trade (Eyster et al., 2015). Naive traders who receive differential information develop different beliefs as they neglect information revealed by trades. When beliefs are sufficiently divergent, they agree to speculate against each other and thus generate excessive trade. By means of a simple simulation described in Appendix A.4, we compute for each treatment the potential number of trades that would occur if participants acting in the role of trader 2 were to trade with each other, at the stated levels of their willingness to buy and sell. We find that simultaneous mechanisms generate significantly more potential for trades than the sequential ones. (The “Low Signal Quality” treatments, whose shares of trades do not differ from each other, are the exception.) This analysis, albeit simplistic, supports the conjecture that naive traders who neglect disagreement in beliefs spawn additional trade.

7 Conclusion

How well traders are able to extract information in markets may depend on the markets’ designs over and above ‘rational’ reasons. Although different but isomorphic trading mechanisms should entail the same outcomes, decisions may vary. Our experiments provide an example where a specific subset of inferences are weak: traders in simultaneous markets, where optimal trading requires Bayesian updating on hypothetical outcomes, do not account for the price’s informativeness. They therefore neglect information revealed by others’ investments. However, when the reasoning is simplified to updating on a single realized event, such ‘cursedness’ is mitigated. Traders are thus more likely to

²⁴This uses a measure for informational efficiency (IE) that is standard in the finance literature (see e.g. Brown and Zhang 1997; De Jong and Rindi 2009): $IE = 1 - \frac{E[\text{Var}[\theta|p_3]]}{V[\theta]}$.

detect covert information while focusing on a single outcome. In this sense, the degree of inference and consequently the quality of informational efficiency interact with market design. Of course, this is only a single setting and despite the numerous robustness checks in the paper we must not presume generalizability. It's a stylized experiment, no more and no less. Subsequent work may address, for example, the largely open research question of price efficiency in sequential trading with more than two consecutive traders.

References

- Y. Amihud, R. Mendelson, and B. Lauterbach. Market microstructure and securities values: Evidence from the Tel Aviv Stock Exchange. *Journal of Financial Economics*, 45:365–390, 1997.
- M. H. Bazerman and W. F. Samuelson. I Won the Auction But Don't Want the Prize. *Journal of Conflict Resolution*, 27(4):618–634, 1983.
- D. P. Brown and Z. M. Zhang. Market Orders and Market Efficiency. *The Journal of Finance*, 52(1):277–308, 1997.
- J. D. Carrillo and T. R. Palfrey. No trade. *Games and Economic Behavior*, 71(1):66–87, jan 2011.
- B. G. Charness and D. Levin. The Origin of the Winner's Curse: A Laboratory Study. *American Economic Journal: Microeconomics*, 1(1):207–236, 2009.
- C. Comerton-Forde, K. Malinova, and A. Park. Regulating Dark Trading: Order Flow Segmentation and Market Quality. 2016.
- V. P. Crawford and N. Iriberry. Level-k auctions. Can a nonequilibrium model of strategic thinking explain the Winner's Curse and overbidding in private-value auctions? *Econometrica*, 75(6):1721–1770, 2007.
- F. De Jong and B. Rindi. *The Microstructure of Financial Markets*. Cambridge edition, 2009.
- I. Esponda and E. Vespa. Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 2014.
- I. Esponda and E. Vespa. Contingent Preferences and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory. 2016.
- J. S. B. T. Evans. *Hypothetical Thinking: Dual Processes in Reasoning and Judgement*. 2007. ISBN 1135419531.
- E. Eyster and M. Rabin. Cursed equilibrium. *Econometrica*, 73(5):1623–1672, 2005.
- E. Eyster, M. Rabin, and D. Vayanos. Financial Markets where Traders Neglect the Informational Content of Prices. *Working Paper*, 2015.
- U. Fischbacher. Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, 2007.
- B. Greiner. An Online Recruitment System for Economic Experiments. In K. Kremer and V. Macho, editors, *Forschung und wissenschaftliches Rechnen, GWDG Bericht 63*, pages 79–93. Gesellschaft für wissenschaftliche Datenverarbeitung, 2004.
- D. M. Grether. Testing bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior and Organization*, 17(1):31–57, 1992.

- S. Grossman. On the Efficiency of Competitive Stock Markets Where Trades Have Diverse Information. *The Journal of Finance*, 31(2), 1976.
- H. Hong and J. C. Stein. A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. *The Journal of Finance*, 54(6): 2143–2184, 1999.
- A. Ivanov, D. Levin, and M. Niederle. Can Relaxation of Beliefs Rationalize the Winner’s Curse?: An Experimental Study. *Econometrica*, 78(4):1435–1452, 2010.
- J. H. Kagel and D. Levin. The winner’s curse and public information. *American Economic Review*, 76:894–920, 1986.
- J. H. Kagel and D. Levin. *Common value auctions and the winner’s curse*. Princeton University Press, 2009.
- J. H. Kagel, D. Levin, R. C. Battalio, and D. J. Meyer. First-price common value auctions: bidder behavior and the ”winner’s curse”. *Economic Inquiry*, 27(2):241–258, 1989. ISSN e.
- D. Kahneman. Maps of Bounded Rationality: Psychology for Behavioral Economics. *The American Economic Review*, 93(5):1449–1475, 2003.
- C. Koch and S. P. Penczynski. The Winner’s Curse : Conditional Reasoning & Belief Formation. 2014.
- A. S. Kyle. Continuous Auctions and Insider Trading. *Econometrica*, 53(6): 1315–1335, 1985.
- D. Levin, J. H. Kagel, and J.-F. Richard. Revenue Effects and Information Processing in English Common Value Auctions. *The American Economic Review*, 86(3):442–460, 1996.
- S. Li. Obviously Strategy-Proof Mechanisms. 2016.
- A. Madhavan. Trading Mechanisms in Securities Markets. *The Journal of Finance*, 47(2):607–641, 1992.
- K. Malinova and A. Park. Liquidity, volume and price efficiency: The impact of order vs. quote driven trading. *Journal of Financial Markets*, 16(1):104–126, 2013.
- G. L. Murphy and B. H. Ross. Predictions from Uncertain Categorizations. *Cognitive psychology*, 27:148–193, 1994.
- T. Odean. Volume, Volatility, Price, and Profit When All Traders Are Above Average. *The Journal of Finance*, LIII(6), 1998.
- R. W. Olshavsky. Task complexity and contingent processing in decision making: A replication and extension. *Organizational Behavior and Human Performance*, 24(3):300–316, 1979.
- M. Pagano and A. Roell. Transparency and Liquidity: A Comparison of Auction and Dealer Markets with Informed Trading. *Journal of Finance*, 51(2):579–611, 1996.

- S. Pouget. Financial market design and bounded rationality: An experiment. *Journal of Financial Markets*, 10(3):287–317, 2007.
- B. H. Ross and G. L. Murphy. Category-Based Predictions : Influence of Uncertainty and Feature Associations. *Journal of experimental psychology. Learning, memory, and cognition*, 22(3):736–753, 1996.
- C. R. Schnitzlein. Call and Continuous Trading Mechanisms Under Asymmetric Information: An Experimental Investigation. *The Journal of Finance*, 51(2): 613–636, 1996.
- E. Shafir and A. Tversky. Thinking through uncertainty: nonconsequential reasoning and choice. *Cognitive psychology*, 24(4):449–474, 1992.
- E. Theissen. Market structure, informational efficiency and liquidity: An experimental comparison of auction and dealer markets. *Journal of Financial Markets*, 3(4):333–363, nov 2000.

A Appendix

A.1 Descriptive Statistics

We compute the share of buys for different ranges of signal values. Table A1 refers to the trades with a transaction price that lies below its prior expectation of 130. The observations in Table A2 refers to rounds with transaction prices above 130. The rows “Diff.” show the differences between the shares in the sequential and simultaneous mechanisms, for the main, the “Low Signal Quality” and the “No Player 1” treatments, respectively.

Table A1: Share of buys at low prices for varying signal intervals

Treatment	All signals	[0 - 0.2]	[0.2 - 0.4]	[0.4 - 0.6]	[0.6 - 0.8]	[0.8 - 1]
SEQ	.4138 (.031)	.0825 (.032)	.16 (.034)	.3704 (.065)	.7647 (.058)	.9048 (.043)
SIM	.4847 (.041)	.12 (.037)	.3053 (.061)	.6818 (.054)	.875 (.048)	.8545 (.065)
Diff.	-.0709	-.0375	-.1453**	-.3114***	-.1103	.0503
<i>N</i>	736	197	170	147	116	118
LSQ-SEQ	.4106 (.044)	.1667 (.042)	.2432 (.060)	.4231 (.067)	.6 (.081)	.6897 (.067)
LSQ-SIM	.5714 (.037)	.3171 (.063)	.4783 (.066)	.56 (.072)	.7407 (.056)	.7733 (.066)
Diff.	-.1606***	-.1504*	-.2351**	-.1369	-.1407	-.0836
<i>N</i>	736	160	143	153	156	133
NP1-SEQ	.3403 (.030)	.0392 (.023)	.1048 (.035)	.3333 (.051)	.6538 (.058)	.8462 (.060)
NP1-SIM	.4495 (.033)	.1214 (.042)	.1939 (.048)	.5789 (.062)	.8088 (.058)	.9245 (.044)
Diff.	-.1092**	-.0823*	-.0891	-.2456***	-.155*	-.0784
<i>N</i>	825	209	203	163	146	118

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ in two-sample t test with unequal variances. CRSE in parentheses.

Table A2: Share of buys at high prices for varying signal intervals

Treatment	All signals	[0 - 0.2]	[0.2 - 0.4]	[0.4 - 0.6]	[0.6 - 0.8]	[0.8 - 1]
SEQ	.6190 (.033)	.0555 (.036)	.2388 (.056)	.4828 (.066)	.9054 (.037)	.9406 (.026)
SIM	.5140 (.033)	- (.)	.1724 (.053)	.2778 (.063)	.7419 (.053)	.84 (.043)
Diff.	.105**	-	.066	.205**	.1635	.1006**
<i>N</i>	692	69	125	130	167	201
LSQ-SEQ	.6151 (.038)	.2537 (.067)	.5294 (.070)	.7 (.066)	.8 (.056)	.7848 (.053)
LSQ-SIM	.3050 (.038)	.2239 (.059)	.2 (.058)	.1818 (.047)	.4464 (.066)	.5079 (.081)
Diff.	.3101***	.0298	.3294***	.5182***	.3536***	.2769***
<i>N</i>	635	134	106	147	106	142
NP1-SEQ	.6738 (.027)	.1475 (.047)	.3889 (.062)	.7 (.063)	.8817 (.042)	.9626 (.018)
NP1-SIM	.4523 (.030)	.1132 (.042)	.0882 (.053)	.225 (.044)	.6813 (.063)	.8302 (.035)
Diff.	.2215***	.0343	.3007***	.475***	.2004***	.1324***
<i>N</i>	821	114	140	170	184	213

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$ in two-sample t test with unequal variances. CRSE in parentheses.

Table A3 shows the shares of buys when prices and signals reflect contrary information because they lie on opposite sides of their corresponding prior expectation. Trading decisions that conform rather with the information in the price than with the information in the signal indicate that participants give thought to the price's informativeness. In all treatment variations, traders 2 in the sequential mechanisms trade more often against the information contained in their own signal: they sell (buy) more often than their peers in the simultaneous mechanism when the price is low (high). The differences between the buys and sells in the two mechanisms are significant for the variations "Low Signal Quality" and "No Player 1".

Table A3: Acting against one's own signal (treatment prices)

	$p_2 \leq 130$ $s_2 > .5$	$p_2 > 130$ $s_2 \leq .5$
SEQ	.7834 (.042)	.2357 (.044)
SIM	.8332 (.036)	.1460 (.040)
Diff.	-.0498	.0877
N	293	277
LSQ-SEQ	.5976 (.059)	.4323 (.047)
LSQ-SIM	.7326 (.049)	.1939 (.044)
Diff.	-.135*	.2383***
N	351	320
NP1-SEQ	.6815 (.049)	.3584 (.040)
NP1-SIM	.8446 (.045)	.1198 (.045)
Diff.	-.1631**	.2386***
N	327	340

Note: *p < 0.1, **p < 0.05, ***p < 0.01 in two-sample t test with unequal variances. CRSE in parentheses.

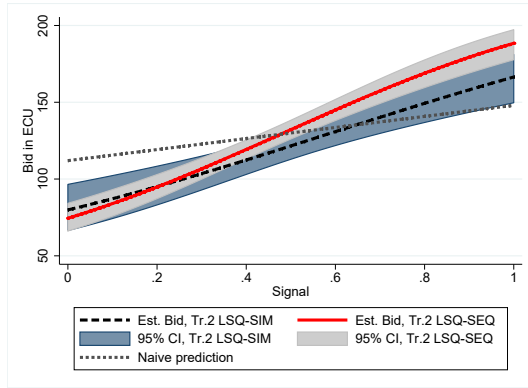


Figure A 1: Estimated average bids in treatments LSQ-SIM and LSQ-SEQ.

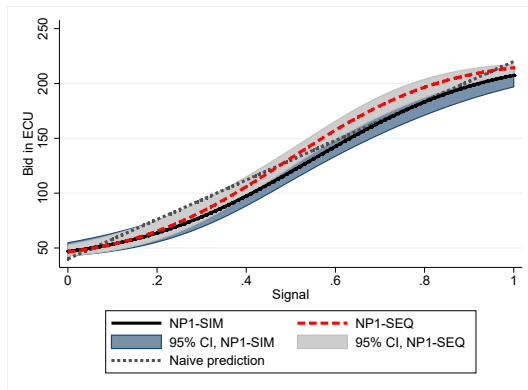
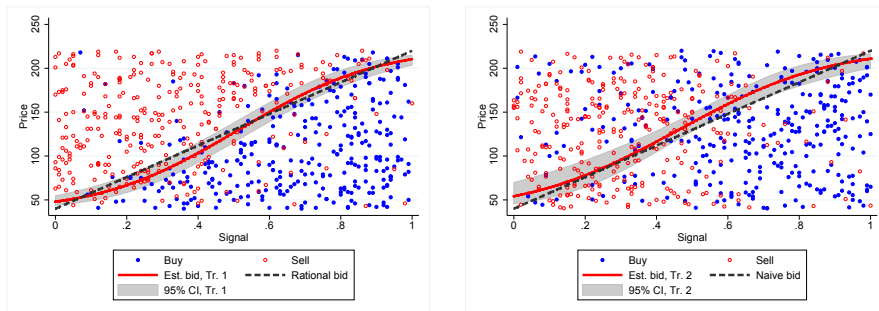


Figure A 2: Estimated average bids in treatments NP1-SIM and NP1-SEQ.



(a) Trader 1

(b) Trader 2

Figure A 3: Buys, sells and estimated average bids of traders 1 (a) and 2 (b) in treatment Hyp-SEQ.

Table A4: Shares of naive decisions

	SIM	SEQ	LSQ-SIM	LSQ-SEQ	Hyp-SEQ	NP1-SIM	NP1-SEQ
η	.3760 (.047)	.1851 (.052)	.4449 (.045)	.2222 (.033)	.2830 (.042)	.4527 (.053)	.1767 (.033)
N	118	108	227	261	106	148	181

Note: CRSE in parentheses. Significant difference at 1% level between SIM & SEQ, between LSQ-SIM & LSQ-SEQ and between NP1-SIM & NP1-SEQ (Wald test). Significant difference in 1-sided Gauss test between Hyp-SEQ and SIM ($p=0.022$), and Hyp-SEQ and SEQ ($p=0.081$).

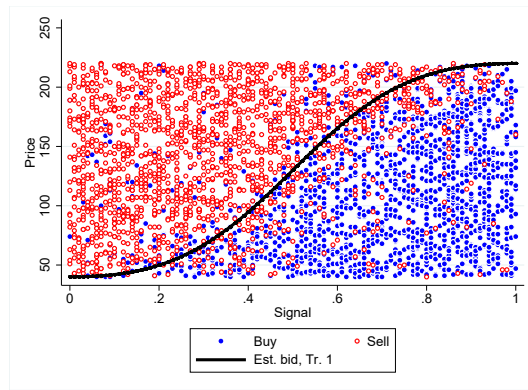


Figure A 4: Bid function for trader 1 given random utility model estimates.

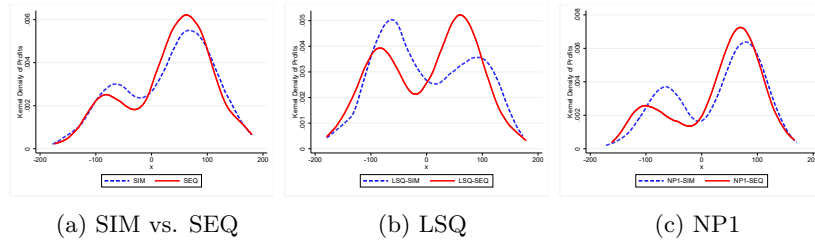


Figure A 5: Kernel density of profits of traders 2 in treatments SIM, SEQ, LSQ-SIM, LSQ-SEQ and NP1-SIM, NP1-SEQ.

Table A5: Profits of traders 2

	Mean	S.E.	Median
SIM	27.63	2.98	44
SEQ	30.65	2.86	43.25
LSQ-SIM	-1.24	3.19	-18.25
LSQ-SEQ	.85	3.21	21
HYP-SEQ*	27.48	4.30	43.25
NP1-SIM	25.30	2.78	50.5
NP1-SEQ	28.36	2.65	52.5

Note: S.E. refers to standard errors of mean. *Excluding rounds that generated zero profit in Hyp-SEQ because no trade occurred.

A.2 Multiple Binomial Testing

This section describes how we identify significant deviations from naive bidding. We test the hypothesis that the propensity to buy conforms with the probability of buying with naive expectations. With naive expectations, the probability that a trader buys equals her posterior belief for the high asset value, which (given uniform priors) equals the signal's value. Thus, the null hypothesis of naive posterior beliefs corresponds to:

$$H_0 : \pi(s^j) = s^j, \quad j = 1, \dots, 9.$$

We round signals to decimals. We merge extreme signals close to 0 and 1 to the nearest category to satisfy testing criteria in the approximate binomial test. We then perform (one-sided) binomial tests for each of the 9 categories. The first column of Table A6 denotes the alternative hypothesis H_A for each test. The alternative hypothesis is chosen to reject naiveté in favor of Bayesian probabilities. The other columns in Table A6 report the p-values for each test for the corresponding treatment.

We account for the multiple testing problem using the Bonferroni significance level of 0.0055 (with a significance level of $\alpha = .05$ for individual tests). Two treatments, SEQ and NP1-SIM, display trading decisions that significantly differ from the naive prediction. In treatment SEQ, the more extreme trading decisions lead to a rejection of the null, while in treatment SIM the share of buys is consistent with naive beliefs. In treatments Hyp-SEQ, NP1-SIM and NP1-SEQ the null is rejected in four out of 9 categories, but only in treatment NP1-SIM the null is rejected after correcting for the multiple testing problem. This significant deviation in the simultaneous mechanism is driven by the overall increased tendency to sell, especially at low signal values. Figure A 2 reveals an estimated bidding curve that lies below the naive function for almost all signal values.

Table A6: P-values in one-sided binomial testing

H_A	SIM	SEQ	Hyp-SEQ	NP1-SIM	NP1-SEQ
$\pi < .1$.7695	.2716	.9999	.9303	.7555
$\pi < .2$.4754	.0364	.3399	.1616	.0194
$\pi < .3$.1751	.0083	.2765	.0002*	.0369
$\pi < .4$.1320	.2614	.0214	.0110	.1658
$\pi \neq .5$.7962	.2642	.7854	.0114	.9146
$\pi > .6$.4000	.0874	.0427	.7092	.0201
$\pi > .7$.1084	.0009*	.0206	.0293	.0808
$\pi > .8$.0506	.1063	.0103	.7250	.0227
$\pi > .9$.9962	.3770	.7435	.9228	.5131

Note: * $p < 0.0055$ (Bonferroni significance level.). Tests for $H_0 : \pi = .5$ are two-sided.

For the treatments with low signal quality, the likelihood for the high asset value is bounded in $[.4, .6]$ due to the signal's low precision. The null adjusts to:

$$H_0 : \pi(s^j) = 0.4 + 0.2 \cdot s^j, \quad j = 1, \dots, 11.$$

Table A7: P-Values in multiple binomial testing

H_A	LSQ-SIM	LSQ-SEQ
$\pi < .4 + .2 * .0$.0422	.3546
$\pi < .4 + .2 * .1$.0191	.0000*
$\pi < .4 + .2 * .2$.0495	.0235
$\pi < .4 + .2 * .3$.0133	.0195
$\pi < .4 + .2 * .4$.0000*	.3714
$\pi \neq .4 + .2 * .5$.4060	.3172
$\pi > .4 + .2 * .6$.4643	.0722
$\pi > .4 + .2 * .7$.0209	.0347
$\pi > .4 + .2 * .8$.6458	.0000*
$\pi > .4 + .2 * .9$.0191	.0007*
$\pi > .4 + .2 * 1$.3047	.5000

Note: * $p < 0.0045$ (Bonferroni significance level.)
Tests for $H_0 : \pi = .5$ are two-sided.

The multiple binomial tests detect in both treatments LSQ-SIM and LSQ-SEQ significant deviations from the share of buys that would be expected under naiveté. The deviations occur at both low and high signal values, reflecting the higher steepness of the bidding curves shown in Figure A 1. The information asymmetry helps trader 2 to take into account the price's informativeness.

A.3 Learning

To investigate whether participants learn over time, we divide observations into two time subsections: an early time interval for the rounds one to ten and a late interval for later rounds. In the subset of price-signal realizations where naive and Bayesian predictions differ, the proportion of naive decisions does not change significantly over time in all treatments except treatment LSQ-SEQ, as shown in Table A8. Furthermore, plotting the share or number of naive decisions across periods does not display any systematic pattern of decay. Even pooling treatments into simultaneous and sequential variants does not reveal any learning effect.

Table A8: Proportion of naive decisions

	SIM	SEQ	LSQ-SIM	LSQ-SEQ	Hyp-SEQ	NP1-SIM	NP1-SEQ
First 10	.3971 (.060)	.2127 (.074)	.4741 (.052)	.2810 (.046)	.3077 (.065)	.5128 (.070)	.1596 (.044)
Last 10	.34 (.079)	.1639 (.058)	.4144 (.068)	.1714 (.038)	.2593 (.057)	.3857 (.073)	.1954 (.044)
Diff.	.0571	.0488	.0597	.1096**	.0484	.1271	-.0358
N	118	108	227	261	106	148	181

Note: *p < 0.1, **p < 0.05, ***p < 0.01. CRSE in parentheses.

A.4 Trading volume

We calculate the number of trades that would occur within one treatment if traders 2 were allowed to trade with each other (as price-takers). To this end, we compare the actual buys and sells that took place at each price values, rounding the latter to the nearest ten. The minimum of buys or sells at a given price value defines the number of transactions that would have been possible between the set of traders 2 at this price. Table A9 shows the share of potential trades per price value, which corresponds to the ratio of potential trades to the maximum possible trading volume. Since every trade requires two trading parties, the maximum number of possible trades at a specific price equals the frequency of this price value divided by two. The simultaneous mechanisms entail significantly more potential trades, except for the treatment variation with “Low Signal Quality” that displays similar shares of trades in each mechanism.

Table A9: Average simulated trading volume with random matching of trader 2 participants

	SIM	SEQ
Main treatments	.8611 (.004)	.7806*** (.004)
Low Signal Quality	.7629 (.006)	.7735 (.007)
No Player 1	.87 (.005)	.6977*** (.003)

***: Share is significantly smaller than in the alternative treatment in a one-sided t-test with $p < .01$.