Learning from unrealized versus realized prices*

M. Kathleen Ngangoué   Georg Weizsäcker

September 2019

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 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

*We benefited from helpful discussions with Dan Benjamin, Erik Eyster, Steffen Huck, Ryan Oprea, Emanuel Vespa, Roberto Weber, as well as with four referees, the Co-Editor (John Asker) and the audiences at St Andrews, Berkeley, DIW Berlin, FU Berlin, MPI Bonn, Cologne, Columbia, CREST, Hamburg, UCL, NYU, Pittsburgh, PSE, UC San Diego, UC Santa Barbara, USC, 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

alia. 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 case in point 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; 2019 and Li 2017) suggest that the task is more accessible in a sequential trading mechanism than in a simultaneous one.

Juxtaposing simultaneous and sequential mechanisms allows two insights. First, it provides a clean identification of naive decisions that are due to the necessity of forming price-contingent strategies. The failure of contingent reasoning is shown to be outcome relevant.[2] Second, it helps assessing how retail investors in real-world markets react to the market structure. Retail investors are more likely to suffer from cognitive biases than institutional investors (see Skiba and Skiba, 2017, for a review) and regulations of retail trading need to assess potential drivers of investors' welfare, including behavioral biases.[3] To the extent that behavior deviates from perfect rationality, a behavioral experiment can complement theoretical considerations in this regard. While a clear distinction between pure simultaneous and sequential markets does not exist in the real world, some market structures have clear features of sequentiality, and others of simultaneity. A prominent example of the latter are order-driven markets, especially those with call auctions, which require investors to supply liquidity without knowledge of liquidity demand (Malinova and Park, 2013; Comerton-Forde et al., 2016). For instance, 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. More generally, and 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.[4]

The stylized dichotomy in trading mechanisms has led to an extensive discussion of the efficiency of simultaneous versus sequential mechanisms. The discussion's emphasis lies, however, on information aggregation, i.e., the reflec-

[1]Experiments 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).

[2]We differentiate between *naiveté*, which is defined in Section 3 as the confusion of conditional with unconditional expectations, and *failure of contingent reasoning*—a possible source of naiveté. Naiveté may more generally be present to some extent, but our treatment exposes the scope of suboptimal inferences caused by a failure of contingent reasoning.

[3]As we discuss below, such an assessment may differentiate between these biases' effect on the investors who suffer from the bias and the resulting mis-pricing effect on other investors.

[4]In line with the difficulty in using price-contingent strategies, Linnainmaa (2010) relates suboptimal trading patterns of retail investors, such as contrarian trades and the disposition effect, to the use of limit orders.

tion of dispersed information in the price, rather than signal extraction, i.e., the inference by traders who observe the price.[5] Importantly, incomplete signal extraction may also influence the speed of price discovery when new information flows into the market. The literature has noted that prices in real and experimental call markets adjust relatively slowly to incoming information (Amihud et al., 1997; Theissen, 2000). As we explain below, our evidence identifies the failure of contingent reasoning as a possible explanation also of this pattern.

Our experimental 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.[6] 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. Contingent reasoning 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 measure the extent to which naiveté fits the data better than the empirical best response, separately by treatment.

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,

[5]Relevant theoretical and experimental studies include Kyle 1985; Madhavan 1992; Pagano and Roell 1996; Copeland and Friedman 1991, 1992; Cason and Friedman 1997, 2008; Schnitzlein 1996; Theissen 2000; Pouget 2007, inter alia. A consensus is that the consolidation of orders allows simultaneous markets like call auctions to aggregate information. Pouget'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.

[6]Traders 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.

a seller receives it. In contrast, in SEQ, where transaction prices are known beforehand, asset demand uses the information contained in the price significantly more, yielding trades that are closer to the empirical best response. Averaging over all situations where the naive prediction differs from the empirically optimal trade, the frequency of naive trading is twice as high in SIM relative to SEQ, at 37% 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, and they may therefore neglect the price as an information source. 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 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. In a separate treatment we draw subjects' attention to one contingency by presenting a single possible price that may realize. Participants submit their hypothetical buy/sell preference at this one possible, but not yet realized, price.[7] The rate of optimal choices in this treatment lies mid-way between that of the two main treatments, illustrating that the difficulty of 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 (Subsection 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. We find that the degree of naiveté 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 demand therefore decreases too fast with the price. The economic bearing of

[7]In the language of Esponda and Vespa (2019), this treatment singles out the contingent preference that may not necessarily coincide with the original preference.

the effect is further discussed in Section 6. We examine the predictions of Hong and Stein (1999) and Eyster et al. (2019) 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. Their erroneous speculation therefore 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 trader therefore learns less from the price. But naiveté is detrimental not only to later players: also the observed payoffs of our market participants 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.[8] Our experiment suggests a specific manifestation of this effect, namely that drawing the attention to the realized price may enable a more rational interpretation of 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). 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 others' actions and behave strategically. Other recent studies directly investigate the importance of contingent reasoning. Subjects appear to systematically disregard relevant information that is conveyed by future, not yet realized events: Charness and Levin (2009) and Koch and Penczynski (2018) show that overbidding in simultaneous mechanisms decreases when finding the optimal solution does not necessitate updating on future events.[9][10] Esponda and Vespa (2019) reduce the extent of anomalies

[8] Shafir 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).

[9] Charness and Levin (2009) analyze the Winner's Curse in a takeover game, whereas Koch and Penczynski (2018) focus on common-value auctions.

[10] Note that inferences from hypothetical events differ from inferences from absence of events

in various decision problems with an experimental protocol that better elucidates payoff consequences of each contingency. The experiments in Martínez-Marquina et al. (2019) point out the role of uncertainty in the failure to engage in contingent thinking. By comparing a stochastic and deterministic variant of the acquiring-a-company game they show that the cognitive load of having to reason through several outcomes does not fully account for the difficulty with contingent reasoning. This difficulty has also given rise to theoretical concepts like 'obviously dominant strategies' (Li, 2017) or 'obvious preferences' (Zhang and Levin, 2017).

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 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.[11] We denote agent i's positive or negative unit demand with $X_i \in \{1, -1\}$. The asset is worth $\theta \in \{\underline{\theta}, \overline{\theta}\}$, with equal probabilities. For all trades, 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 p_i, and vice versa if the asset is sold:

$$\Pi_i = (\theta - p_i)X_i \tag{1}$$

Traders do not observe the fundamental value θ but they each receive a private signal $s_i \in [0, 1]$.[12] The true value θ determines which of two triangular densities the signal is drawn from, such that in the low-value state participants receive low signals with a higher probability, and vice versa:

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

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

like in Jin et al. (2018) and Brown et al. (2012). In our setting, the absence or presence of events itself does not reveal any additional information because strategies are contingent on realized events only, even in simultaneous mechanisms where events are realized at a later point in time.

[11]Because 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.

[12]A binary signal structure would have been easier to implement in an experiment but would also have reduced the identifiability of naive trader 2 behavior. With a binary signal and an efficient pricing rule, a naive trader differs from a rational trader only in that he is indifferent between buying and selling in cases where both traders receive identical signals. With continuously-valued signals, more differences between naive and rational trading arise.

As we are interested in informational inefficiencies in markets, we consider an environment with a natural source of strategic uncertainty: trader 1's decision determines the informativeness of trader 2's price. Trader 1 and 2 therefore face separate transaction prices p_i. We first describe the (uninformative) price for trader 1 as well as his trading task, which is identical in both treatments. We then describe the resulting (informative) price for trader 2 and her task, which differs across treatments.

2.1 Task of trader 1

Trader 1's price p_1 is uniformly distributed in $[\underline{\theta}, \overline{\theta}]$ and is uninformative about the fundamental value θ. Trader 1 receives a private signal s_1, as described above, and states his maximum willingness to pay by placing a limit order b_1. If p_1, which is unknown at the time of b_1's submission, lies weakly below b_1, the trader buys one unit of the asset. If p_1 strictly exceeds b_1, he sells one unit.[13]

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 Subsection 2.2.2.) Let Z_1 denote the indicator function that takes on value 1 if a limit order is reversed, and 0 otherwise. Trader 1's demand is $X_1(p_1, b_1)$:

$$X_1(p_1, b_1) = Y_1(p_1, b_1)(1 - Z_1) - Y_1(p_1, b_1)Z_1 \tag{3}$$

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

The demand resulting from a standard limit order corresponds to $Y_1(p_1, b_1)$, while $X_1(p_1, b_1)$ incorporates the trader's choice of reversing the limit order. We will denote demand by X_i instead of $X_i(p_i, b_i)$ whenever it is unambiguous.

In sum, trader 1's demand is based on a single piece of information, the signal s_1, and is expressed by a single bid, b_1. As we will describe in Section 3.1, the optimal bid b_1 increases linearly in s_1.

2.2 Task of trader 2

Trader 2 is informed that the price p_2 reflects the expectation of an external market maker who observes trader 1's buying or selling decision and, on that basis, infers information about trader 1's signal s_1. Importantly, to avoid any ambiguity in the description, trader 2 learns the pricing rule that maps p_1 and the realized value of X_1 into p_2:

$$p_2 = \begin{cases} \frac{\overline{\theta}+p_1}{2}, & \text{if } X_1 = 1 \\ \\ \frac{\underline{\theta}+p_1}{2}, & \text{if } X_1 = -1 \end{cases} \tag{4}$$

[13] The design does not allow for a "no trade" option that 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.

This mapping reflects the Bayesian inference on behalf of the market maker, assuming rational bidding of trader 1.

Along with the equation for the price, participants 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. The rule implies that p_2 is informative about the asset value θ: trader 1's private signal s_1 influences p_2 through X_1. It is therefore optimal for trader 2 to condition her investment decision on both s_2 and p_2.

In the two main treatments, trader 2 faces the same pricing rule but different decision tasks.

2.2.1 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, b_2. Importantly, and different from the sequential treatment, trader 2 does not know her own price p_2 when submitting b_2. She buys a unit of the asset if $p_2 \leq b_2$, and otherwise she sells a unit of the asset. Optionally, she can change her bid into a reversed limit order.

2.2.2 Sequential treatment (SEQ)

In treatment SEQ, trader 2 observes the price p_2 as specified in (4) *before* making her decision. The game proceeds sequentially, with trader 1 first choosing his bid b_1. As in treatment SIM, his demand X_1 determines the price for trader 2, p_2. Trader 2 observes the realized values of $\{p_1, p_2, s_2\}$ and then chooses between buying and selling at p_2.

It is straightforward to check that treatments SIM and SEQ are strategically equivalent. For any given s_2, treatment SEQ allows for four possible strategies contingent on $p_2 \in \{\frac{\underline{\theta}+p_1}{2}, \frac{\overline{\theta}+p_1}{2}\}$: $\{buy, buy\}$, $\{buy, sell\}$, $\{sell, buy\}$ 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. This implies that strategy spaces are equivalent between treatments SIM and SEQ, and equally so for the corresponding payoff consequences. Rational responses to a fixed belief about trader 1 would therefore lead to the same purchases and sales in the two treatments.[14]

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

[14] This statement holds under the assumptions of subjective utility theory. Generalizations of expected utility can lead to different choices across treatments. For instance, while in treatment SEQ buying/selling choices at known prices can be described as choices over two simple lotteries, setting limit orders in SIM corresponds to choices between three compound lotteries with uncertainty over dichotomous asset values and prices. Hence, violations of the compound independence axiom or of the reduction of compound lotteries could also generate inconsistencies across treatments (cf. Karni and Safra, 1987; Segal, 1988, 1990; Keller et al., 1993). Our experiment was not designed to compare between non-expected utility theories.

rational expectations and properly updates on her complete information set. As the third benchmark prediction, we consider the case that trader 2 fully neglects the price's informativeness. For all three predictions, we assume traders to be risk-neutral.[15]

3.1 Rational best response of trader 1

The optimal bid function of trader 1 is straightforward. Trader 1 has only his private signal s_1 to condition on. 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 (3)) that b_1^* increases linearly in the signal:

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

3.2 Rational best response of trader 2

Under rational expectations about trader 1's strategy, trader 2 maximizes her expected payoff via 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:[16]

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

The optimal bidding of trader 2 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.

Even in the current design where prices can only take two possible values, the exercise of inferring the correct information from the pricing rule in (4) remains complex. Using this additional information does, however, not necessarily require extreme sophistication as the Bayes-Nash (BN) strategy of trader 2 simplifies here to a step function: p_2 reflects the market maker's expectation (see (4)), 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. Notice that the extreme prediction in form of a step function results directly from the efficient pricing rule p_2, which equals correct expectations under Bayes-Nash. If the price p_2 is only a noisy reflection of market expectations, then the Bayesian strategy becomes a less steep, S-shaped curve.

In this sense, 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 equal the exact expectation if trader 1 deviates from $X_1(p_1, b_1^*)$. For a stronger test of naive beliefs, we consider the empirical best response (EBR)

[15]Risk aversion shifts the bid function toward the more cautious trading strategy of buying below (and selling above) the *ex-ante* expected price, which avoids large losses and gains. Risk aversion alone, however, cannot explain differences between treatments.

[16]For 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 would be forgone.

to the participants acting as trader 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, pooled across treatments SIM and SEQ, for the calculation of EBR), together with the naive prediction that we describe next.[17] The graphs represent the 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 sufficiently positive information (high s_2).

3.3 Best response to naive beliefs of trader 2

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:

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

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. (2019) that best responds to the belief that agent 1's equilibrium mixture over bids arises regardless of their information.[18] Note that while naiveté can be modeled with level-k reasoning or fully cursed beliefs, these concepts do not distinguish between sequential and simultaneous decisions and do, per se, not account for differences between treatments. Similarly, Li (2017)'s concept of obviously dominant strategies would not account for differences between simultaneous and sequential treatments because none of the strategies in treatment SEQ is obviously dominant. For any price-contingent buy and sell, the asset's stochastic value may generate payoffs in both the gain and loss domain. Li's general idea may nevertheless apply in that it may be easier to compare relevant payoffs conditional on a single price.

[17] The kinks in the EBR function arise because of the numerical approximation to the fixed point. For close approximation, signals are rounded to lie on a grid with step size 0.1.

[18] In 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 $F(\frac{b-\underline{\theta}}{\overline{\theta}-\underline{\theta}}) = F(s_1)$, with density $\frac{1}{2}f(s_1|\underline{\theta}) + \frac{1}{2}f(s_1|\overline{\theta}) = 1$. The analysis of Eyster and Rabin (2005) and Eyster et al. (2019) also allows for intermediate levels of naiveté, 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.

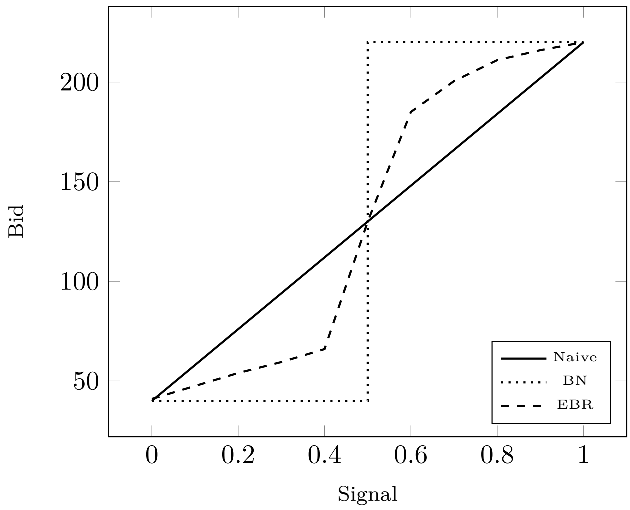

Figure 1: Naive, Bayes Nash and empirical best responses of trader 2.

3.4 Hypotheses

As outlined in the Introduction, we conjecture that 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.

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

Participants read the instructions for both roles, traders 1 and 2, before learning which role they were assigned to. Instructions included an elaborate computer-based simulation of the signal structure as well as a comprehension test. The support of the asset value was $\{40, 220\}$.[19] Each session lasted approximately 90 minutes and participants earned on average EUR 22.02. Total earnings consisted of a show-up fee of EUR 5.00, an endowment of EUR 15.00 and profits from the randomly drawn round (which could be negative but could not deplete the entire endowment). Units of experimental currency were 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, reversed limit orders (3% 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 consider the set of trades $X_2(s_2, p_2)$ with prices and signals between the solid (b_2^{Naive}) and the dashed (b_2^{EBR}) bidding functions: $\{X_2(s_2, p_2) \in \{-1, +1\} \mid \big((s_2 < 0.5) \cap (b_2^{EBR} \leq p_2 \leq b_2^{Naive})\big) \cup \big((s_2 > 0.5) \cap (b_2^{Naive} \leq p_2 \leq b_2^{EBR})\big)\}$. 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. Within this relevant area, we calculate the proportion

[19] See the Online Appendix for a set of instructions for treatments SIM and SEQ. We chose the mildly unusual asset values $\{\underline{\theta}, \overline{\theta}\} = \{40, 220\}$ in an attempt to avoid midpoint effects.

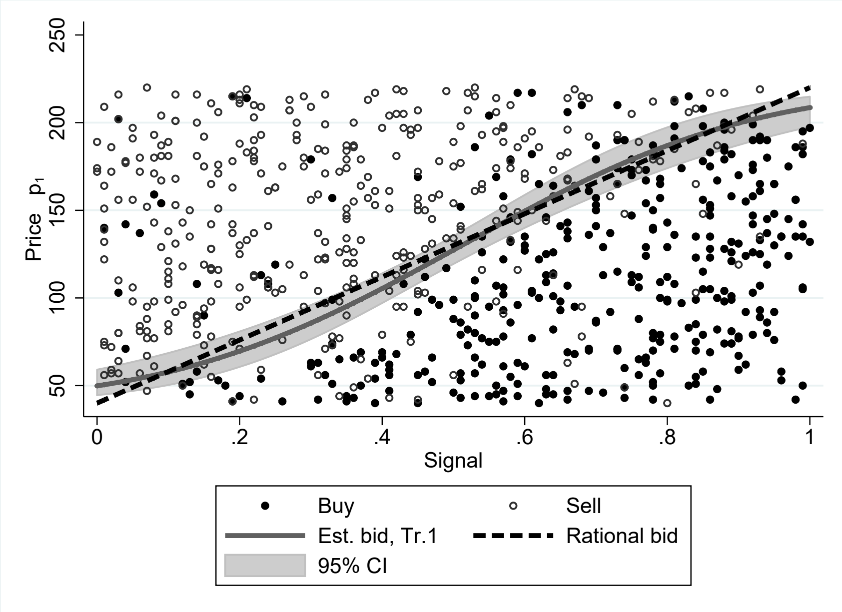

Note: The average bidding curve corresponds to $\underline{\theta} + (\overline{\theta} - \underline{\theta}) \cdot \hat{P}(X_1|s_1)$, where $\hat{P}(X_1|s_1)$ is the probit estimate of the probability of trader 1 buying in treatment SIM.

Figure 2: Trades of traders 1.

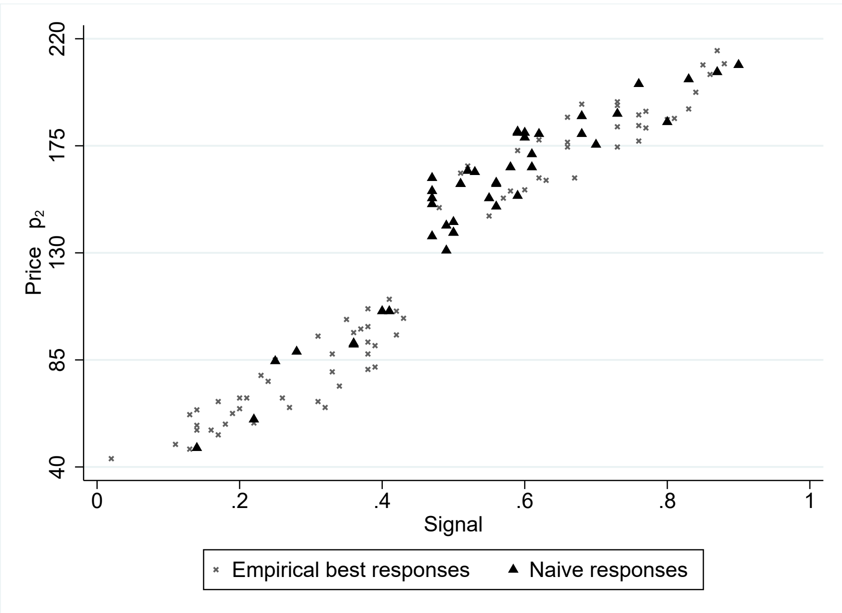

Figure 3: Sells and buys of trader 2 within the relevant area in treatment SIM.

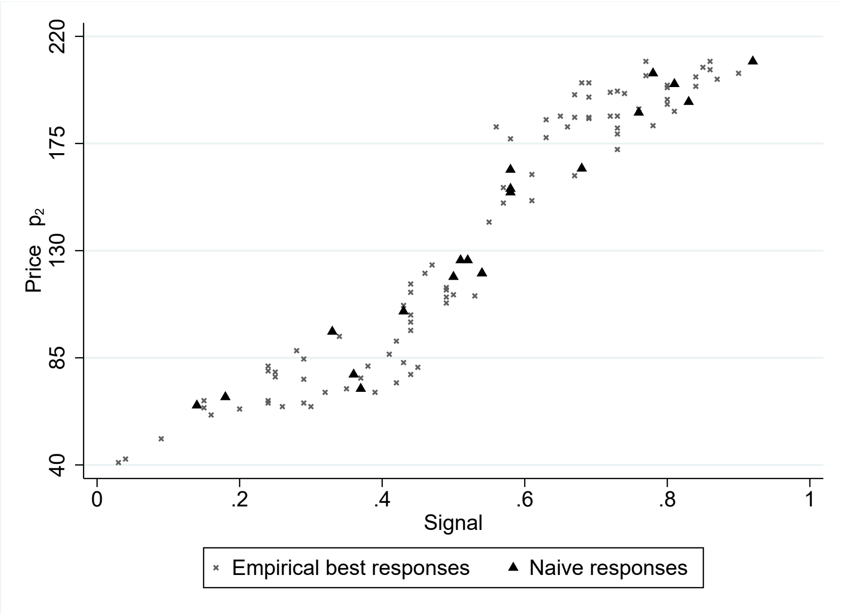

Figure 4: Sells and buys of trader 2 within the relevant area in treatment SEQ.

η of naive decisions:

$$\eta = \frac{d_N}{d_N + d_B} \tag{8}$$

where d_N and d_B denote the number of orders consistent with naive (triangle markers in Figures) and EBR (cross markers) predictions, 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 A2 shows that the share of naive decisions in treatment SIM ($\eta = 0.37$) is twice as large as in treatment SEQ ($\eta = 0.19$).[20] The difference is statistically significant ($p = 0.0091$, Wald test).[21]

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. 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. 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.

The same conclusion is reached with an alternative measure that uses the complete set of observations. We estimate the mean bid function and compute

[20] Appendix Table A2 shows that this effect is robust for high and low signals (i.e., $s_2 > 0.5$ vs. $s_2 \leq 0.5$) but it is generally larger for high signals.

[21] All presented tests account for robust standard errors clustered at the subject level.

its Euclidean distance to the EBR benchmark (as shown in Figure 1). We find that in both treatments average bidding functions significantly deviate from the EBR benchmark ($p < 0.001$ in Kolmogorov-Smirnov test for both treatments) but subjects in treatment SEQ bid, on average, closer to the EBR than their peers in treatment SIM (mean distance of 13.74 ECU in SEQ vs. 22.02 ECU in SIM). This is true for all signals (see Appendix Figure A 1).

In Appendix A.3, we also consider the evolution of decisions in the course of the experiment. Participants tend to bid less naively over time in most treatments but these experience effects are mostly non-significant.[22]

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 perceivably, 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.

This evidence of trader 2's full naiveté nevertheless requires a cautious interpretation. Outside the relevant area, the use of reversed limit orders draws a differentiated picture. While reversing the limit order is a strictly suboptimal strategy for trader 1, the Bayes-Nash best response of trader 2 can also be sustained with extreme reversed limit orders. Most empirical best responses cannot be made with reversed limit orders, but their use may hint at trader 2's partial sophistication. In line with this, subjects in the role of trader 2 reverse their limit order more often than those in the role of trader 1 (15% for trader 2 vs. 3% for trader 1 in treatment SIM). Dynamics shed additional light on whether these reversed limit orders might be due to confusion or partial sophistication: subjects in the role of trader 1 use fewer reversed limit orders over time, suggesting that they recognize reversing as a dominated strategy, whereas those in the role of trader 2 significantly increase their use of reversed limit orders over time (e.g. an increase of 11% to 19% in SIM between early versus late rounds, see Appendix A.3). Yet, in the aggregate, these reversed limit orders do not generate more Bayes-rational trades.

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.

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 than in the main treatments. The densities in the new treatments are depicted

[22] Random effects probit regressions with a reciprocal time trend do, however, detect significant experience effects in treatments SIM and LSQ-SEQ, as we present in the Online Appendix.

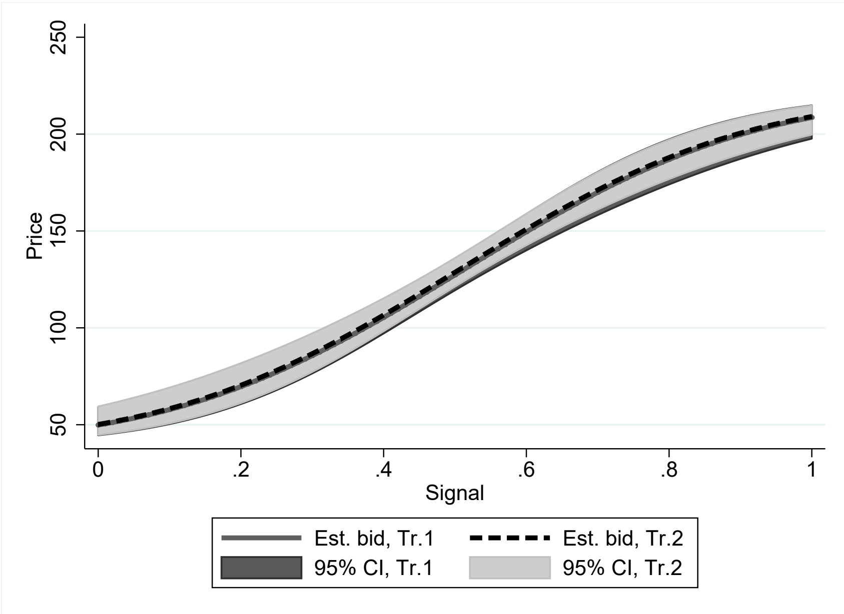

Figure 5: Estimated average bids of traders 1 and 2 in treatment SIM.

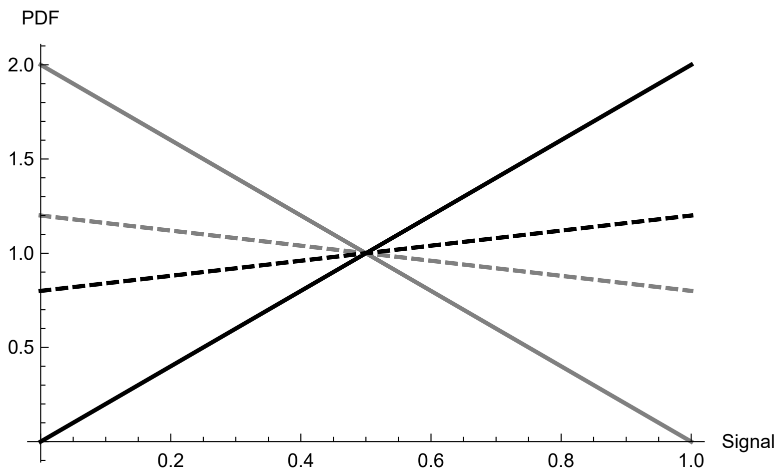

Figure 6: Signal distributions for trader 1 (solid) and trader 2 (dashed) in LSQ treatments.

in Figure 6 and take the following form.

$$f(s_i|\theta = \underline{\theta}) = 1 - \tau_i(2s_i - 1)$$
$$f(s_i|\theta = \overline{\theta}) = 1 + \tau_i(2s_i - 1)$$
$$\texttt{with } \tau_1 = 1 \texttt{ and } \tau_2 = 0.2.$$

The Bayes-Nash prediction remains unchanged relative to the main treatments and the naive prediction adjusts by generating a "flatter" reaction to s_2. The empirical best response is mildly "steeper" in the LSQ treatments, as it is optimal to follow the price more. Our analysis accounts for this difference.

Behavior of trader 2 deviates from the naive prediction in both treatments LSQ-SIM and LSQ-SEQ, as participants indeed show a steeper response (see Figure A 3). From the probit estimation alone it is, however, not clear whether this steeper response comes from an increased sensitivity to price or signal. The estimation of a random utility model in the Online Appendix suggests that participants in treatment LSQ-SIM consider the information contained in prices even less, compared to treatment SIM.

Regarding the robustness of the main result, we observe that 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, different from LSQ-SEQ at $p = 0.0003$, Wald test). Also, participants in the role of trader 2 of LSQ-SEQ act more frequently against their own signal (see Appendix Table A1). In sum, the importance of trading mechanisms for rational decision making prevails under the new informational conditions.

5.2 Strategic uncertainty

Strategic uncertainty makes for part of the game's complexity. 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 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 labelled "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 when observing the demand $X_1(p_1, b_1)$ of a trader 1 who behaves rationally.[23]

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

[23]The distributions of the additional signals (one for each asset value) are shown in a graphical display. The instructions do not explain the distributions' origins.

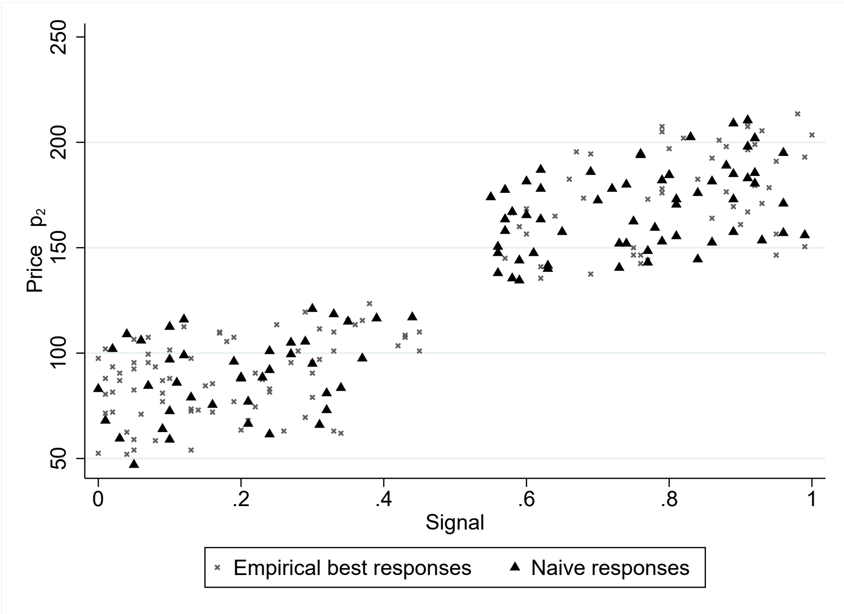

(a) LSQ-SIM

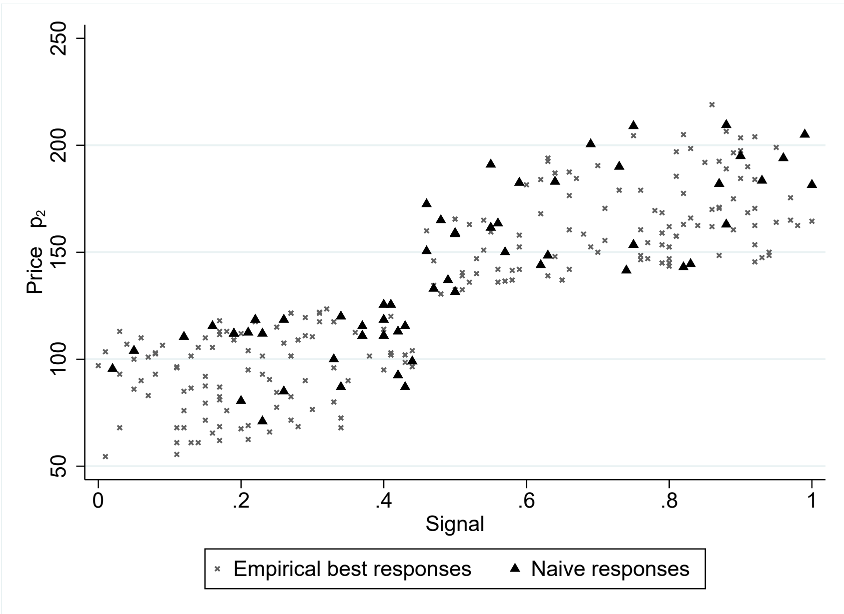

(b) LSQ-SEQ

Figure 7: Trader 2's buys and sells consistent with either naive bidding or EBR in LSQ-SIM and LSQ-SEQ, respectively.

signal that is correlated with the asset's value. Like in the main treatments, the instructions display the updating rule (4) 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 effect of simultaneous versus sequential trading persists. The share of naive decisions is two and a half times higher in NP1-SIM than 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. 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 stem from the additional cognitive strain that conditional thinking requires. Perhaps, it is not conditionality per se that is difficult for the participants, but the fact that they have to make two decisions in treatment SIM (one for each possible price realization) and 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 for a single price. However, they do so conditionally, for a single candidate price $\hat{p}_2$ that is equiprobably drawn from the two price values that are possible after updating via rule (4). 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 contingent, 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.[24]

Figure 9 and Appendix Table A2 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 (0.37 versus 0.28, p=0.077, 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=0.071, one-sided t test).[25] Altogether, we conclude from the above tests that reducing the number of hypothetical trading decisions

[24]Esponda and Vespa (2019) account for irrational decisions by allowing their so-called original, contingent and conditional preferences to be different primitives. In our experiments, treatments SIM, Hyp-SEQ and SEQ elicit the original, contingent and conditional preferences, respectively.

[25]Notice that the lower rate of suboptimal decisions in Hyp-SEQ relative to SIM is consistent with the main idea of Li's (2017) obvious strategy proofness: in Hyp-SEQ, the set of relevant prices is reduced to a singleton, helping the participants to detect the optimal strategy.

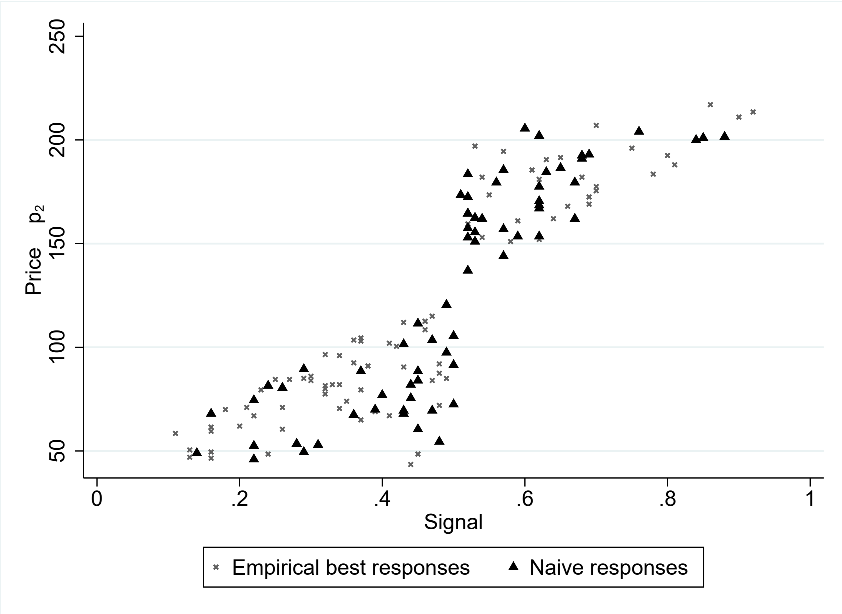

(a) NP1-SIM

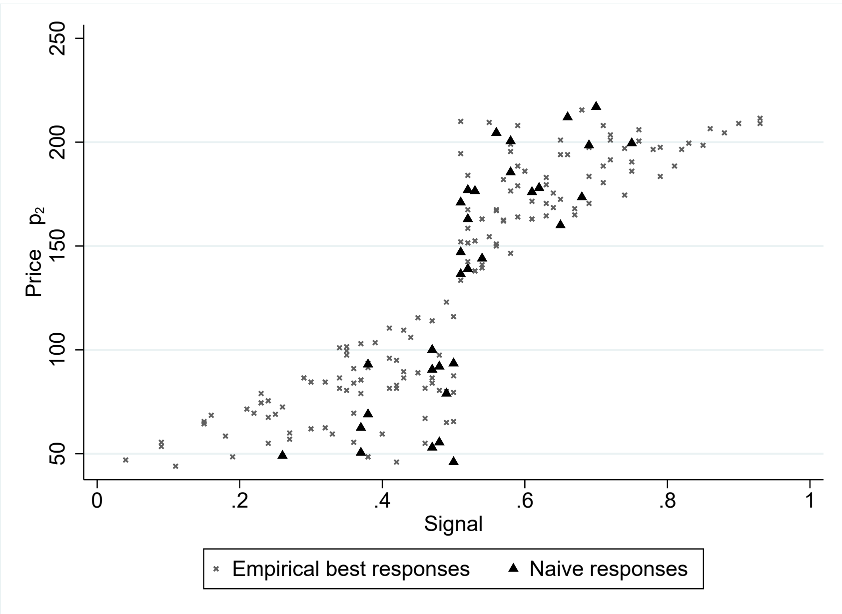

(b) NP1-SEQ

Figure 8: Trader 2's buys and sells consistent with either naive or Bayesian bidding in NP1-SIM and NP1-SEQ, respectively.

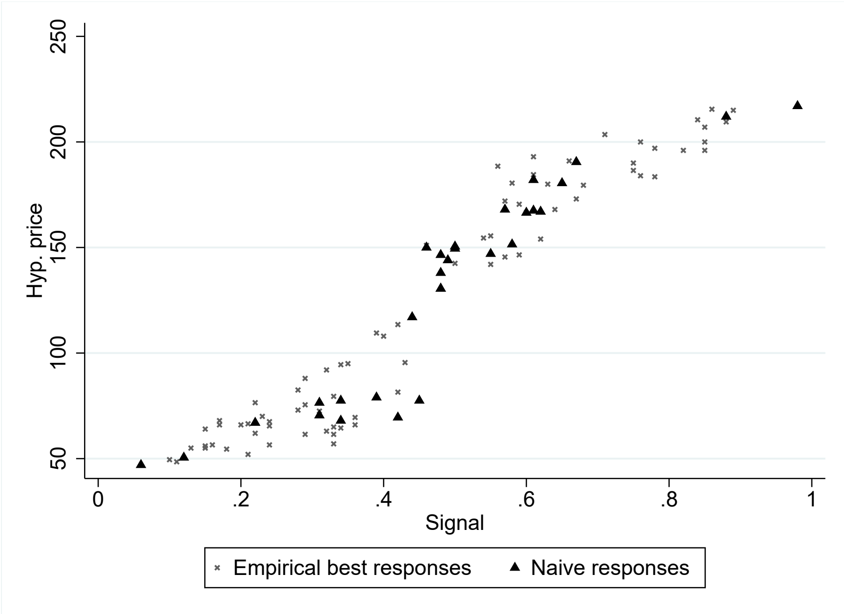

Figure 9: Trader 2's buys and sells consistent with either naive bidding or EBR in Hyp-SEQ.

reduces the degree of naiveté, but does not eliminate it.[26] [27]

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 here but estimates for individual treatments can be found in Appendix Table A4.) 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:

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

[26]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 2015 experiment, however, also has other differences to the present one.

[27]This finding is consistent with the experimental results in Martínez-Marquina et al. (2019), showing that the difficulty with contingent reasoning cannot be reduced to the difficulty of thinking through several outcomes.

with

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

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

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 subject-specific 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 would correctly weight the signal, $\beta = 1$, but would ignore the information in the price, $\alpha = 0$. An intermediate level of naiveté 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^{***}	2.54^{**}	1.36^{***}
	(0.31)	(0.90)	(0.36)
α	-	0.60^{*}	1.85^{***}
		(0.26)	(0.22)
λ	0.0314^{***}	0.0230^{***}	0.0373^{***}
	(0.003)	(0.004)	(0.006)
σ_u	0.0010	0.0010	0.0039
N	3435	2220	2260

Note: $^{*}p < 0.10, ^{**}p < 0.05, ^{***}p < 0.01$. Maximum simulated likelihood estimation simulates subject-specific random effects assuming $u_i \sim N(0, \sigma_u)$. 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 7).[28] 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 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 ($\alpha = 1.85$).[29] While inferences in the sequential mechanisms are also not optimal (with the scale parameter λ capturing the disturbances), the coefficients α and β in simultaneous and sequential mechanisms significantly differ from each other ($p<0.001$ in Chow likelihood ratio test with and without heterogeneous scale parameters λ). 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 on 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.

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 Hong and Stein (1999) and Eyster et al. (2019). 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

[28]Overweighting one's own signal also appears in some, but not all, of the related inference problems in herding experiments. See e.g. Goeree et al. (2007), and the meta analysis in Weizsäcker (2010).

[29]This 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.

towards its ex-ante mean.[30]

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 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:[31]

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

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, both traders either buy or sell in 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. If, in contrast, trader 2 bids naively, then she tends 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 than in treatment SIM.

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).

Another way to assess price efficiency in the two treatments is to compare the variance of fundamental value conditional on the price, $Var[\theta|p_3]$. 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 ($p < 0.001$, nonparametric median test, taking each market as a unit of observation) and with a somewhat sizeable difference: in treatment SIM, the

[30] Hong 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. The model in Eyster et al. (2019) uses partially cursed equilibrium to study mispricing, using a more standard (and more static) model of financial markets with incomplete information akin to that in Grossman (1976).

[31] In 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.

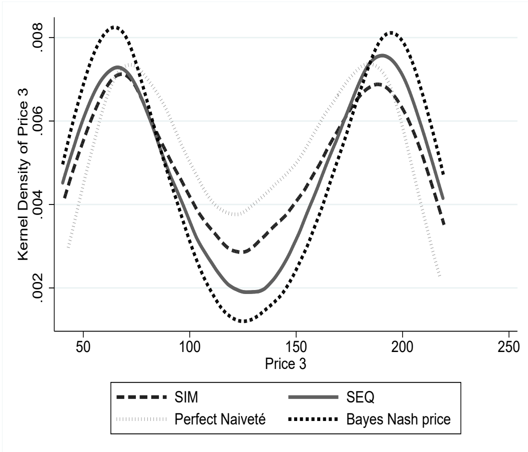

(a) s_1, s_2 in same direction

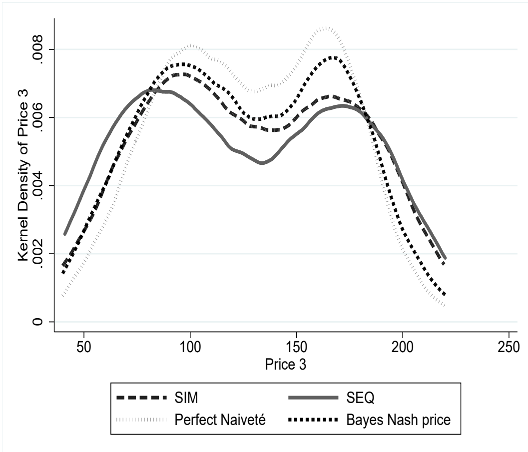

(b) s_1, s_2 in opposite direction

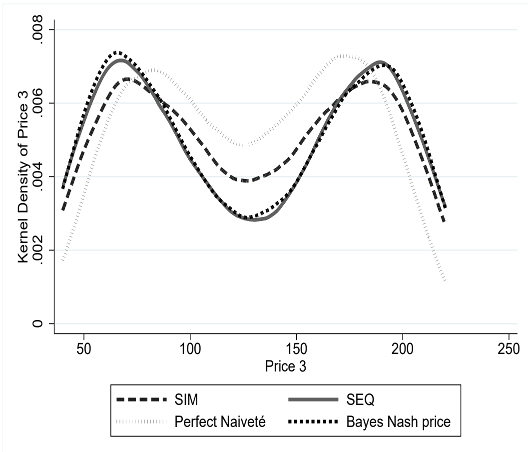

(c) All signals

Figure 10: Kernel density of efficient price p_3 after naive, rational and actual demand of traders 1 and 2 in SIM and SEQ. The computation of price p_3 is based on the empirical distribution of the price p_2.

price explains on average 21% of the variance in the asset value, versus 27% in treatment SEQ.[32] In addition, in our experiment naive bidding leads to weakly smaller profits and could have generated increased trading volume (see Appendix A.1 and A.4).

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 in theory, 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 naiveté is mitigated. Traders are thus more likely to 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.

[32] This 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): $IE = 1 - \frac{E[Var[\theta|p_3]]}{V[\theta]}$.

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.

B. A. L. Brown, C. F. Camerer, and D. Lovallo. To Review or Not to Review ? Limited Strategic Thinking at the Movie Box Office. *American Economic Journal: Microeconomics*, 4(2):1–26, 2012.

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. Y. T. N. Cason and D. Friedman. Price Formation in Single Call Markets. *Econometrica*, 65(2):311–345, 1997.

T. N. Cason and D. Friedman. A Comparison of Market Institutions. In *Handbook of Experimental Economics Results*, volume 1. 2008.

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.

T. E. Copeland and D. Friedman. Partial Revelation of Information in Experimental Asset Markets. *The Journal of Finance*, XLVI(1):265–295, 1991.

T. E. Copeland and D. Friedman. The Market Value of Information : Some Experimental Results. *The Journal of Business*, 65(2):241–266, 1992.

V. P. Crawford and N. Iriberri. 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. *Working Paper*, 2019.

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. *The Journal of Finance*, 74(1):371–399, 2019.

U. Fischbacher. Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171–178, 2007.

J. K. Goeree, T. R. Palfrey, B. W. Rogers, and R. D. McKelvey. Self-correcting information cascades. *Review of Economic Studies*, 74(3):733–762, 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.

G. Z. Jin, M. Luca, and D. Martin. Is no news (perceived as) bad news? An experimental investigation of information disclosure. 2018.

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.

D. Kahneman. Maps of Bounded Rationality: Psychology for Behavioral Economics. *The American Economic Review*, 93(5):1449–1475, 2003.

B. Y. E. D. I. Karni and Z. V. I. Safra. "Preference Reversal" and the Observability of Preferences by Experimental Methods. *Econometrica*, 55(3): 675–685, 1987.

R. L. Keller, U. Segal, and T. Wang. The Becker-DeGroot-Marschak Mechanism and Generalized Utility Theories: Theoretical Predictions and Empirical Observations. *Theory and Decision*, 34:83–97, 1993.

C. Koch and S. P. Penczynski. The winner's curse: Conditional reasoning and belief formation. *Journal of Economic Theory*, (174):57–102, 2018.

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. *American Economic Review*, 107 (11):3257–3287, 2017.

J. T. Linnainmaa. Do Limit Orders Alter Inferences about Investor Performance and Behavior ? *The Journal of Finance*, LXV(4):1473–1506, 2010.

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.

A. Martínez-Marquina, M. Niederle, and E. Vespa. Failures in Contingent Reasoning : The Role of Uncertainty. 2019.

G. L. Murphy and B. H. Ross. Predictions from Uncertain Categorizations. *Cognitive psychology*, 27:148–193, 1994.

K. Ngangoue and G. Weizsäcker. Learning from unrealized versus realized prices. 2015.

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.

U. Segal. Does the Preference Reversal Phenomenon Necessarily Contradict the Independence Axiom ? *The American Economic Review*, 78(1):233–236, 1988.

U. Segal. Two-Stage Lotteries without the Reduction Axiom. *Econometrica*, 58 (2):349–377, 1990.

E. Shafir and A. Tversky. Thinking through uncertainty: nonconsequential reasoning and choice. *Cognitive psychology*, 24(4):449–474, 1992.

A. Skiba and H. Skiba. Institutional Investors. In H. K. Baker, G. Filbeck, and V. Ricciardi, editors, *Financial Behavior: Players, Services, Products, and Markets*. Oxford University Press, New York, 2017.

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.

G. Weizsäcker. Do We Follow Others when We Should ? A Simple Test of Rational Expectations. *The American Economic Review*, 100(5):2340–2360, 2010.

L. Zhang and D. Levin. Bounded Rationality and Robust Mechanism Design : An Axiomatic Approach. *American Economic Review: Papers & Proceedings*, 107(5):235–239, 2017.

A Online Appendix

A.1 Descriptive statistics and additional figures

Trading against private information. Table A1 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 gave thought to the price's informativeness. In all treatment variations, traders 2 in the sequential mechanisms traded more often against the information contained in their own signal: they sold (bought) more often than their peers in the simultaneous mechanism when the price was low (high). Differences between buys and sells in the two mechanisms are significant for the variations "Low Signal Quality" and "No Player 1".

Table A1: Acting against one's own signal: share of buys when signal and price move in opposite directions

	$p_2 \leq 130$ $s_2 > .5$	$p_2 > 130$ $s_2 \leq .5$
SEQ	.7834	.2357
	(.042)	(.044)
SIM	.8332	.1460
	(.036)	(.040)
Diff.	-.0498	.0877
N_{SEQ}/N_{SIM}	157/136	140/137
LSQ-SEQ	.5976	.4323
	(.059)	(.047)
LSQ-SIM	.7326	.1939
	(.049)	(.044)
Diff.	-.135*	.2383***
N_{SEQ}/N_{SIM}	164/187	155/165
NP1-SEQ	.6815	.3584
	(.049)	(.040)
NP1-SIM	.8446	.1198
	(.045)	(.045)
Diff.	-.1631**	.2386***
N_{SEQ}/N_{SIM}	148/179	173/167

Note: $^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$ in two-sample t test with unequal variances. CRSE in parentheses. Numbers of observations N depend on the number of signal price combinations (s_2, p_2) within a category and therefore differ.

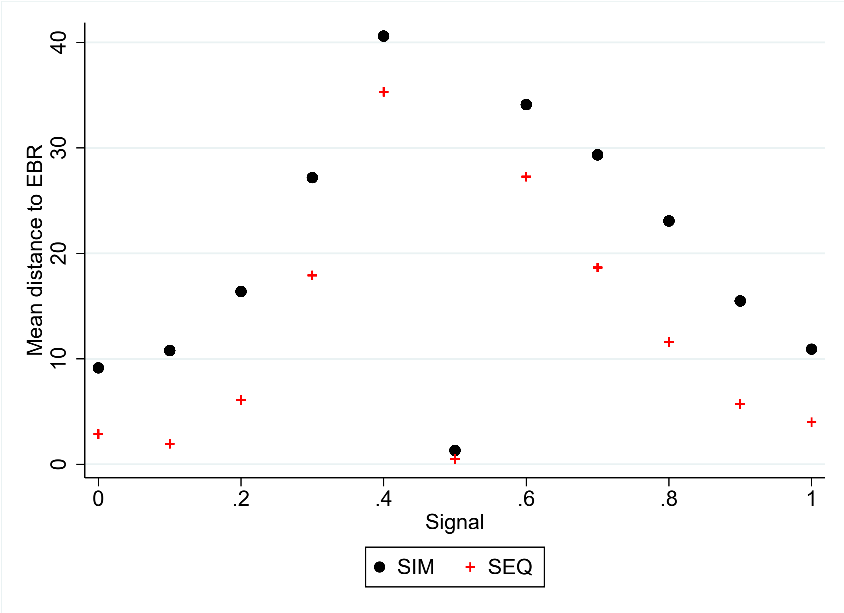

Figure A 1: Euclidean distance between estimated bid function of trader 2 and EBR in SIM and SEQ.

Figure A 1 shows the Euclidean distance between the estimated bid function of each treatment and the joint EBR function computed using the pooled sample of participants in the role of trader 1 in SIM and SEQ and signals grouped to bins of 0.1. As expected, the distance is smaller for extreme and uninformative signals ($s_i \approx 0, s_i \approx 0.5, s_i \approx 1$) than for other signals. More importantly, the distance is always smaller in treatment SEQ than in treatment SIM, showing that participants in SEQ bid closer to the EBR.

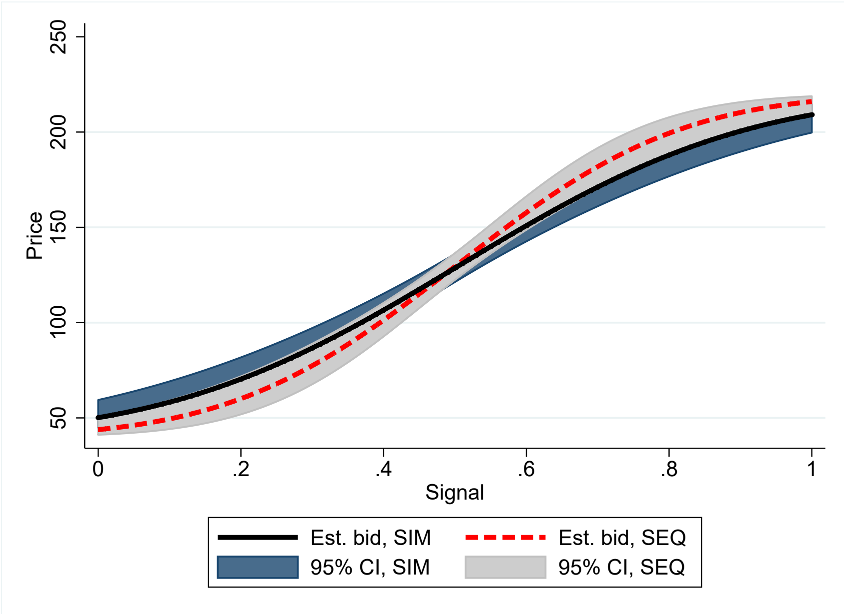

Figure A 2: Estimated average bids of trader 2 in treatments SIM and SEQ.

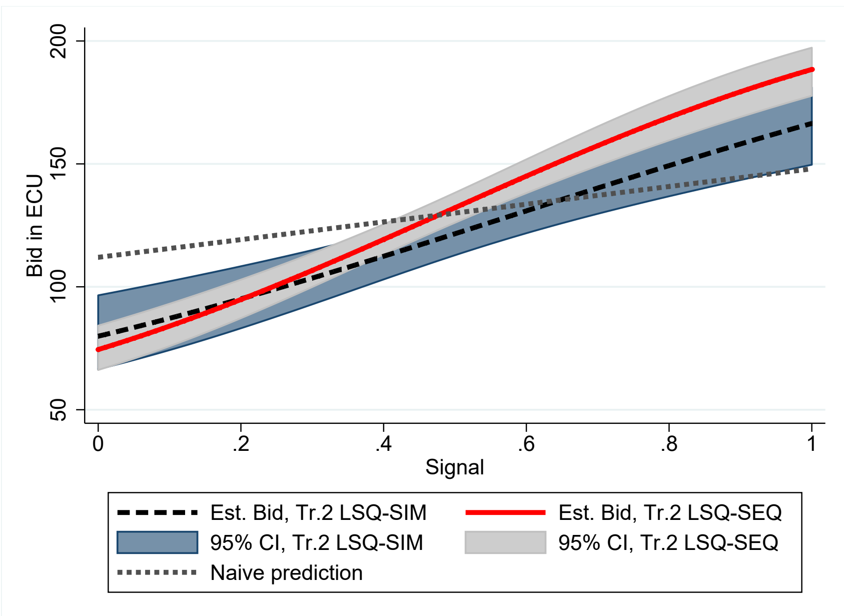

Figure A 3: Estimated average bids of trader 2 in treatments LSQ-SIM and LSQ-SEQ.

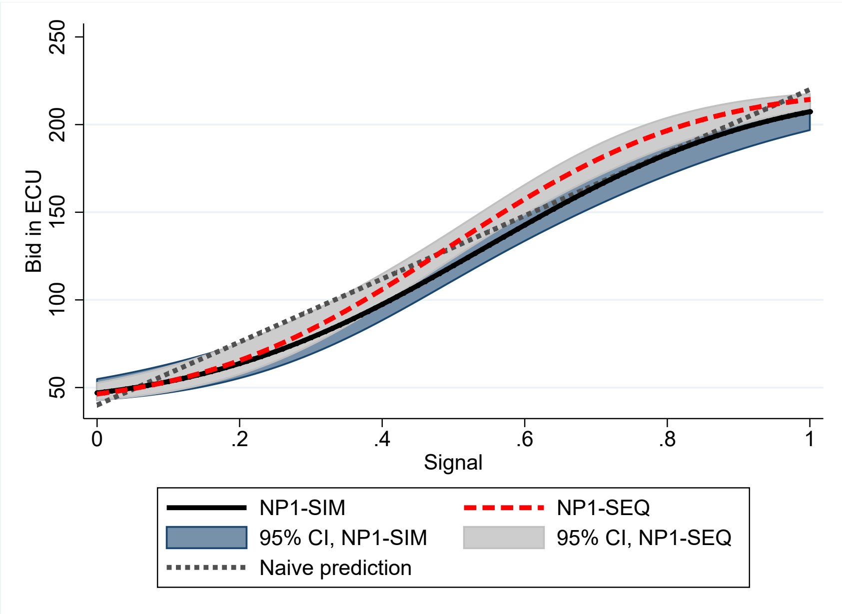

Figure A 4: Estimated average bids of trader 2 in treatments NP1-SIM and NP1-SEQ.

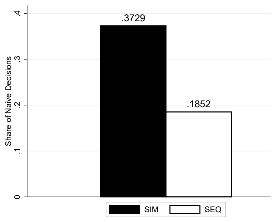

(a) Main

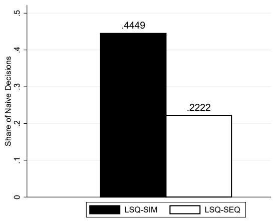

(b) Low Signal Quality

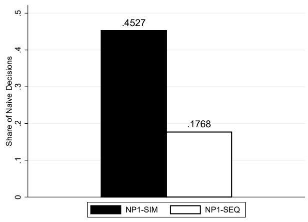

(c) No Player 1

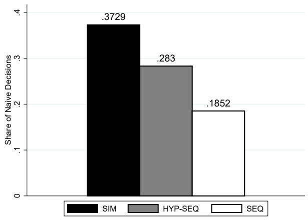

(d) Main vs HYP-SEQ

Figure A 5: Share of naive decisions across treatment variations

Table A2: Share of naive decisions (η) for all, high and low signals

	All s_2	$s_2 \leq 0.5$	$s_2 > 0.5$
SIM	0.373	0.300	0.448
	(0.047)	(0.073)	(0.065)
SEQ	0.185	0.137	0.228
	(0.052)	(0.045)	(0.064)
Difference	0.188***	0.163**	0.220***
N_{SIM}/N_{SEQ}	118/108	60/51	58/57
LSQ-SIM	0.445	0.35	0.545
	(0.045)	(0.049)	(0.066)
LSQ-SEQ	0.222	0.254	0.187
	(0.033)	(0.051)	(0.038)
Difference	0.223***	0.096*	0.358***
$N_{LSQ-SIM}/N_{LSQ-SEQ}$	227/261	117/138	123/110
NP1-SIM	0.453	0.376	0.556
	(0.053)	(0.064)	(0.082)
NP1-SEQ	0.177	0.148	0.204
	(.033)	(0.041)	(0.052)
Difference	0.276***	0.228***	0.352***
$N_{NP1-SIM}/N_{NP1-SEQ}$	148/181	85/88	63/93
Hyp-SEQ	0.283	0.307	0.25
	(0.042)	(0. 072)	(0.047)
Difference to SEQ	-0.098*	-0.17**	-0.022
Difference to SIM	0.09*	-0.007	0.198***
$N_{Hyp-SEQ}$	106	62	44

Note: CRSE in parentheses. One-sided t tests with ***: p<0.01; **: p<0.05; *: p<0.10. Numbers of observations N depend on the number of signal price combinations (s_2, p_2) within the relevant area and therefore vary.

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. Tables A3 shows mean and median profits of each treatment. 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 6b): 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.

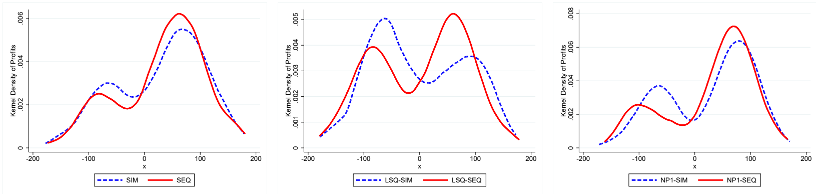

(a) SIM vs. SEQ (b) LSQ (c) NP1

Figure A 6: Kernel density of profits of traders 2 in treatments SIM, SEQ, LSQ-SIM, LSQ-SEQ and NP1-SIM,NP1-SEQ.

Table A3: 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 Random utility model by treatment

Table shows the results of the maximum likelihood estimation by treatment variation.

Table A4: Results of MSL estimation

	SIM	SEQ	LSQ-SIM†	LSQ-SEQ	NP1-SIM	NP1-SEQ	Hyp-SEQ
β	1.487	1.267	21.388*	5.406*	2.335	0.482**	6.856
	(0.82)	(0.44)	(12.41)	(3.31)	(2.42)	(0.26)	(7.16)
α	0.839	1.692**	≈ 0	2.273**	0.261*	1.726**	0.615
	(0.43)	(0.34)	(0.002)	(0.76)	(0.56)	(0.32)	(1.52)
λ	0.0293***	0.0422***	0.0089***	0.0238**	0.0245***	0.727**	0.012***
	(0.01)	(0.01)	(0.001)	(0.01)	(0.007)	(0.03)	(0.003)
σ_u	0.0142	0.0067	56.67***	0.0011	0.0112	0.004	0.016
	(6.13)	(4.20)	(15.88)	(1.94)	(7.13)	(2.47)	(17.03)
N	720	720	700	680	800	860	620

Note: * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Std. Err. in parentheses. Hypothesis testing for β and α refers to one-sided tests of deviations from 1. †The coefficients for LSQ-SIM are obtained using a constrained optimization with $\alpha, \beta \geq 0$ since an unconstrained optimization results in technically invalid values of $\alpha = -6.10$ (and $\beta = 82.49$).

Across all treatments, α, the weight given to price information, is always considerably lower in the simultaneous than in the sequential mechanism: In the sequential markets participants put too much weight on the price (α is significantly larger than 1) whereas in the simultaneous treatments subjects underweight the informational content of prices. Differences are larger in the robustness treatments LSQ & NP1; in particular in treatment LSQ-SIM where subjects assign extremely high and zero weights to the signal and price, respectively. Treatment SIM exhibits reasonable estimates that do not differ from 1. However, we add for completeness that these estimates change substantially if we estimate the model without the reversed limit orders ($\hat{\alpha} \approx 0$ (0.00009), $\hat{\beta} = 2.22$ (0.311) , std. err. in parentheses).

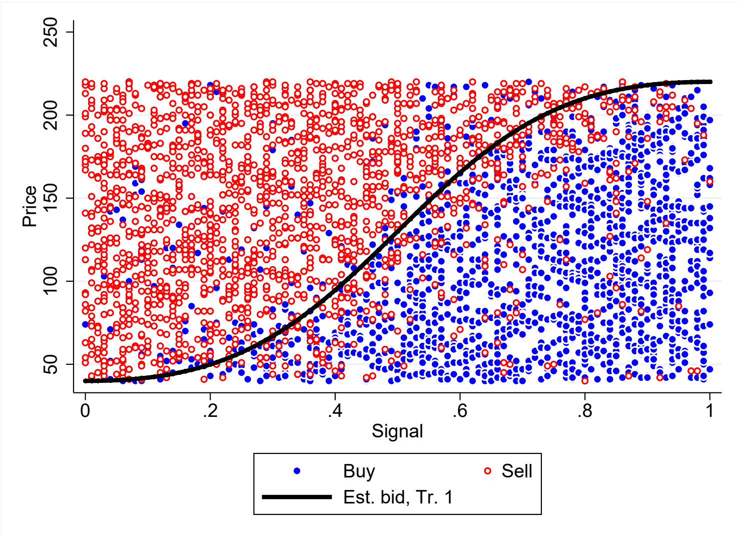

Figure A 7: Bid function for trader 1 given random utility model estimates.

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. We first checked that there are no significant differences in traders 1 trading behavior over time, which implies that the empirical best response of trader 2 is stable. For trader 2, we find only mild evidence that the sequential variant of the game facilitates learning about the other agents' private information. 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 A5. Furthermore, plotting the share or number of naive decisions across rounds does not display any systematic pattern of decay. Even pooling treatments into simultaneous and sequential variants does not reveal significant learning. A random-effects probit regression with reciprocal time trend detects some experience effects in treatment SIM where subjects bid less naively over time. Another type of experience effect can be detected in treatment NP1-SEQ where subjects trade less aggressively after incurring a loss.

Table A5: Proportion of naive decisions

	SIM	SEQ	LSQ-SIM	LSQ-SEQ	Hyp-SEQ	NP1-SIM	NP1-SEQ
First 10	.3971	.2127	.4741	.2810	.3077	.5128	.1596
	(.060)	(.074)	(.052)	(.046)	(.065)	(.070)	(.044)
Last 10	.34	.1639	.4144	.1714	.2593	.3857	.1954
	(.079)	(.058)	(.068)	(.038)	(.057)	(.073)	(.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. Number of observation N depends on the number of signal price combinations (s_2, p_2) that realize within the relevant area where predictions differ and therefore varies.

Table A6: Use of reversed limit orders over time, trader 1

Trader 1	SIM	SEQ	LSQ-SIM	LSQ-SEQ	HYP-SEQ
Round	0.047	0.089	0.089	0.038	0.119
1 - 10	(0.018)	(0.024)	(0.032)	(0.017)	(0.033)
Round	0.017	0.064	0.06	0.024	0.094
11 - 20	(0.010)	(0.026)	(0.030)	(0.013)	(0.036)
Difference	0.030*	0.025	0.029	0.014	0.025

***: p<0.01; **: p<0.05; *: p<0.10. Clustered Robust Standard Errors in parentheses.

Table A7: Use of reversed limit orders over time, trader 2

Trader 2	SIM	LSQ-SIM	NP1-SIM
Round	0.114	0.134	0.078
1 - 10	(0.032)	(0.033)	(0.025)
Round	0.189	0.194	0.115
11 - 20	(0.052)	(0.047)	(0.038)
Difference	-0.075*	-0.060*	-0.037**

***: p<0.01; **: p<0.05; *: p<0.10. Clustered Robust Standard Errors in parentheses.

A.4 Trading volume

Naive beliefs may not only affect prices and profits, but may also trigger speculative trade (Eyster et al., 2019). 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.

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 number of potential transactions that could have been observed at a given price is given by the minimum of buys or sells at this price. The number of potential trades is then normalized by the maximum number of trades. Since every trade requires two trading parties, the maximum number of possible trades at a price equals the frequency of this price value divided by two. Table A8 shows the share of potential trades, which corresponds to the ratio of potential trades to the maximum possible trading volume. 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. This simulation, albeit simplistic, supports the conjecture that naive traders who neglect disagreement in beliefs spawn additional trade.

Table A8: Average simulated trading volume with random matching of trader 2 participants

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

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