

1487

Discussion Papers

Learning from Unrealized versus Realized Prices

Opinions expressed in this paper are those of the author(s) and do not necessarily reflect views of the institute.

IMPRESSUM

© DIW Berlin, 2015

DIW Berlin
German Institute for Economic Research
Mohrenstr. 58
10117 Berlin

Tel. +49 (30) 897 89-0
Fax +49 (30) 897 89-200
<http://www.diw.de>

ISSN electronic edition 1619-4535

Papers can be downloaded free of charge from the DIW Berlin website:
<http://www.diw.de/discussionpapers>

Discussion Papers of DIW Berlin are indexed in RePEc and SSRN:
<http://ideas.repec.org/s/diw/diwwpp.html>
<http://www.ssrn.com/link/DIW-Berlin-German-Inst-Econ-Res.html>

Learning from unrealized versus realized prices^{*}

Kathleen Ngangoue Georg Weizsäcker

June 2015

Abstract

Our market experiment investigates the extent to which traders learn from the price, differentiating between situations where orders are submitted before versus after the price has realized. When market participants have to submit their bids conditional on the price, they show a bias by reacting only to their private information and not to the hypothetical value of the price. In a sequential trading mechanism, where the price is known at the bid submission, bids react to price to an extent that is roughly consistent with the benchmark theory.

JEL-Classification: D82, D81, C91

Keywords: Naive Expectations, Asymmetric Information, Rational Expectations, Sequential Markets

^{*}We benefited from helpful discussions with Dan Benjamin, Erik Eyster, Steffen Huck, Ryan Oprea, Roberto Weber, as well as with the audiences at DIW Berlin and the Berlin Behavioral Economics Workshops. We also thank colleagues at Technical University Berlin and WZB for support in the conduct of the experiment. Financial support by the ERC (Starting Grant 263412) is gratefully acknowledged. Ngangoue: DIW Berlin, Weizsaecker: Humboldt-Universität zu Berlin and DIW Berlin. Corresponding contact: weizsaecker@hu-berlin.de

1 Introduction

Market prices reflect much information about an asset's fundamental value. The extent to which traders are able to utilize this information is difficult to measure as one often lacks control of the traders' restrictions, beliefs and preferences. One possibility to detect an incomplete degree of 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. An important dimension of variability between different markets lies in the conditionality of price. In *simultaneous* markets, the trading price is unknown to the traders at the time when they make their decisions--examples are financial markets with supply/demand function regimes, or simultaneous auctions. Traders should certainly learn from the price, i.e. their bids should optimally condition on it. 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 is consistent with this conjecture, as is the psychological evidence on accessibility (Kahneman, 2003) and contingent thinking (Evans, 2007).¹ 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 the existing evidence renders it plausible that the price is more accessible in a sequential trading mechanism than in a simultaneous one.

Our experiment focuses on markets with diverse information about a risky common-value asset. To trade optimally, an investor 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 simultaneous versus sequential versions. We implement the simultaneous trading game in two different treatments: first as a limit order market (treatment LO), where participants receive a private signal and state their maximum willingness to pay. If the price lies above the limit, the trader buys one unit of the asset, otherwise she sells one unit. (In each treatment, we restrict the trades to a single unit of supply or demand per trader.) At the time they make their decisions, participants in treatment LO do not know the price but would optimally take into account that each possible market price reveals some information about the asset's value. Optimal learning from the price therefore softens the demand's downward reaction to a higher price, relative to the case that the price is

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).

uninformative. The same is true in the game's second simultaneous variant, the price-list treatment (treatment PL): participants receive a list of possible prices and state to each price whether they would buy or sell. Here, too, the information included in the price would reduce a rational trader's demand elasticity, but traders may fail to appreciate this because the price is hypothetical. In contrast, in the sequential market (treatment S), the price is known when the traders make their decisions to buy or sell. Conditional thinking is not necessary here, and we conjecture that this enables a more efficient use of the information contained in the price. Despite the difference in informational conditions, treatments S and PL are economically equivalent in that the traders can make their bids contingent on the price: they have isomorphic strategy sets and isomorphic mappings from strategies to payoffs.² Section 2 presents the experimental design and Section 3 discusses behavioral hypotheses in detail.

The data analysis of Section 4 shows that the participants' ability to infer information from the price varies substantially between simultaneous versus sequential market settings. In simultaneous settings, participants act as if they use their private information only. Price seems to matter only in its direct influence on the utility from trade as is indicated by a comparison between the groups of traders who have uninformative versus informative prices. The analysis also shows that the participants heavily rely on their own signals and even over-weight their signals' importance relative to the theoretical benchmark of risk-neutral Bayesian traders; equivalently, they underweight the prior distribution of asset values. In contrast, in the sequential treatment S, in which transaction prices are known beforehand, asset demand is affected by the information contained in the price. Estimating the strength of taking the information into account (Section 4.3) shows that in the sequential treatment S both pieces of information receive about the same decision weight.

Altogether, the experiment thereby provides evidence of an interaction between market microstructure and the efficiency of information usage. Expressed 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: when a given price is not yet realized, traders behave as if they ignore the connection between other traders' information and the price. Aggregate demand therefore decreases too fast with the price. Across the repetitions of the games we also observe a differential tendency to learn between the different treatments. Learning appears only in the sequential treatment and the difference in trading activity between simultaneous and sequential treatments therefore in-

²Treatment LO differs from them by imposing that the buy/sell decisions are monotonic in price. Therefore, the only 'clean' sequential/simultaneous variation is that between treatment S and treatment PL. Treatment LO is, however, important for the comparison with the case of uninformed prices, as we will discuss in Section 2.

creases over time.

The experimental literatures in economics and psychology provide several sets of related evidence that inference is more complex in simultaneous formats of interaction. Anticipating the informational content of each possible outcome in a simultaneous trading setting involves a substantial amount of mental work. Psychologists have confirmed quite generally that decision processes depend on task complexity (Olshavsky, 1979). In more complex environments, participants prefer approaches demanding less cognitive strain. They focus on one model, one alternative or one relevant category (at a time) 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). One possibility to reduce the complexity of learning or updating tasks is to proceed in a sequential mechanism, like in quote-driven markets.³ 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 processed information more easily and exhibited less non-Nash behavior when facing a take-it-or-leave-it price instead of bidding in a double auction. Similarly, auction experiments 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 substantially as soon as finding the optimal solution does not necessitate updating on future events (Charness and Levin, 2009; Koch and Penczynski, 2014).⁴ 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.

We complement the described evidence on contingent thinking in strategic situations (bilateral bargaining games, auctions and strategic voting games) by addressing a financial market setting where markets clear exogenously. Our sole

³Shafir 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).

⁴Charness and Levin (2009) analyze the Winner's Curse in a takeover game, whereas Koch and Penczynski (2014) focus on common-value auctions.

focus is on the information contained in the price, which the traders may or may not use rationally. The simple structure of the traders' decision problems also helps us to straightforwardly assess whether the typical trader makes too much or too little inference from the price.

2 Experimental Design

Our basic framework is identical across treatments. The experiment simulates a market with a single risky asset and money. A market consists of two agents, agent 1 and agent 2. Both agents either buy or sell one unit of the risky asset.⁵ The asset is worth $\theta = \underline{\theta}$ or $\theta = \bar{\theta}$, each occurring with equal probability. The agents cannot observe the fundamental value θ but they each receive an informative private signal $s_i \in [0, 1]$ about it. 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:

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

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

Each agent i faces a separate transaction price p_i . The price p_1 for agent 1 is uniformly distributed between \underline{p} and \bar{p} and is uninformative about the fundamental value θ . Agent 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. His demand is summarized by the function X_1 :

$$X_1 = \begin{cases} +1 & \text{if } p_1 \leq b_1 \\ -1 & \text{if } p_1 > b_1 \end{cases} \quad \text{where } p_1 \sim U(\underline{p}, \bar{p}) \quad \text{and} \quad (\underline{\theta}, \bar{\theta}) \subset (\underline{p}, \bar{p}) \quad (2)$$

The task of agent 2 varies across treatments. We start by describing the two treatments whose comparison lies at the core of interest, the simultaneous treatment PL and the sequential treatment S.

2.1 Treatment PL

Agent 2 faces a list of possible price values p_2 lying between \underline{p} and \bar{p} . To each of these possible prices she states a decision, buy or sell. The sole payoff-relevant

⁵Because of a possible reluctance to sell short, we avoided any notion of short sales in the experimental instructions. Participants were told that they already possess a portfolio that needs to be adjusted by selling or buying one unit of a specific asset.

decision is the one relating to the actual price p_2 , which she does not know at the time of submitting her decisions. She only knows that p_2 is a weighted average of investor 1's limit order b_1 and a noise term ε :

$$p_2 = \gamma b_1 + (1 - \gamma)\varepsilon \quad (3)$$

Due to its dependence on b_1 , p_2 is influenced by agent 1's private signal s_1 and is therefore informative about the asset's fundamental value. Hence, agent 2 can condition each of her decisions on two different sources of information, s_2 and p_2 .

2.2 Treatment S

In treatment S, agent 2 observes the market price as specified in (3) *before* making her decision. The game proceeds sequentially, with agent 1 first choosing his limit order b_1 based on his private signal. As in the other treatments, his order b_1 partly determines the market price p_2 for agent 2. Agent 2 observes the realized price p_2 in addition to her private signal s_2 before choosing between buying and selling at this price. Agent 2 makes only one decision.

2.3 Treatment LO

Treatment LO also prescribes simultaneous actions, like treatment PL, but asks for a limit order also by agent 2 and hence has the same action space for both agents. They each receive a private signal and simultaneously submit their limit orders before their relevant transaction price realizes. The only difference between both agents remains in the price: while agent 1 encounters an uninformative price p_1 , the price p_2 is again given by (3), i.e. is informative about θ .

The following table summarizes the treatments' different information conditions and action spaces.

Table 1: Treatment variations

	Simultaneous		Sequential
	Limit Order (LO)	Price List (PL)	(S)
Transaction price known	-	-	✓
Action space of agent 2	limit order b_2	{buy, sell} $\forall p_2$	{buy, sell}

2.4 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 agent $i \in \{1, 2\}$ is the difference between the fundamental value and the market price, and vice versa if the asset is sold:

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

$$\text{with } X_i = \begin{cases} +1 & \text{if } i \text{ buys} \\ -1 & \text{else} \end{cases} \quad (5)$$

Note that treatment LO cannot be directly compared with treatment S. By placing a limit order, participants in treatment LO are constrained to monotonic strategies: they buy at low prices and sell at higher prices. Participants of treatment S, however, may sell at high prices, buy at medium-valued prices and sell again when the price is too low. Decisions of treatment S are therefore best compared with treatment PL that also allows for non-monotonic strategies. Notice also that the strategy spaces are isomorphic between treatments PL and S, specifying a buy/sell decision for each p_2 , and that the payoffs arising from each combination of strategies and signals are identical between treatments PL and S. Under any rational response to a fixed belief about agent 1, the traders's decision should lead to the same purchases and sales in these two treatments.

3 Predictions

We compare the participants' behavior to two theoretical predictions. The first benchmark is the case where agent 2 has rational expectations and properly updates on her complete information set. As a second benchmark, we focus on a special form of bounded rationality, namely the case where agent 2 fully neglects the price's informativeness. In both cases, we assume the agents to be risk neutral.

3.1 Bayes (Nash) Best Response

Agent 1 has only his private signal s_1 to condition his bid upon. His optimal limit order b_1^* maximizes the expected profit conditional on s_1 and it is easy to show (using the demand function (2)) that b_1^* increases linearly in the signal:

$$b_1^*(s_1) = \arg \max_{b_1} E[(\theta - p_1)X_1 | s_1] \quad (6)$$

$$= E[\theta | s_1] = \underline{\theta} + (\bar{\theta} - \underline{\theta})s_1 \quad (7)$$

Under rational expectations about agent 1's strategy, agent 2 maximizes her expected payoff conditioning on both her private signal s_2 and the informative market price p_2 . Her maximization problem leads to the following fixed point solution:

$$b_2^*(s_2, p_2) = \arg \max_{b_2} E[(\theta - p_2)X_2 | s_2, p_2] \quad (8)$$

$$= E[\theta | s_2, p_2 = b_2] \quad (9)$$

The optimal bidding behavior of agent 2 typically follows--for a wide range of signals--a cutoff-strategy that switches from buying to selling as the price increases.⁶ The cutoff b_2^* corresponds to the market price that makes her indifferent between a buy and a sell, given both her private signal and the market price. Agent 2's optimal strategy does not linearly increase in s_2 : since both signals s_1 and s_2 are drawn from the same conditional distribution, the correlation between the market price p_2 and the signal s_1 causes agent 2 to react more aggressively to her own signal. A low market price indicates that a low asset value is likely, inducing agent 2 to sell at a lower price than agent 1 (holding their private signals constant). Analogously, the expected value of the asset given both a high signal and a high price exceeds the expectation conditional on only a high signal. The theoretical Bayes Nash (BN) bid thus follows agent 1's own signal in an S-shaped way. Agent 1 reacts strongly to his signal and therefore reacts to price changes with a relatively low probability.

However, the theoretical BN best response is not the only relevant 'rational' benchmark. In the experiment, participants in the role of agent 1 plausibly deviate from their theoretical best response b_1^* and participants in group 2 would optimally adjust to the actual behavior of group 1. We therefore also consider the empirical best response (EBR) to the participants acting as agents 1. The computation of the empirical best response is similar to the fixed point computation of the theoretical BN best response, with the only difference being that it is based on the observed bidding behavior of agent 1, and computed via a numerical approximation to the fixed point.

Both benchmarks BN and EBR are depicted in Figure 1, for the parameters of the actual experiment that are reported in Section 4.

3.2 Best Response to Naive Beliefs

In contrast to optimal behavior, an agent 2 with naive expectations does not infer any information from the price. She fails to account for the connection between

⁶For a set of relatively uninformative signals around 0.5, the expected fundamental value covaries tightly with the price, such that multiple fixed points may exist and the best response may not be a monotonic strategy.

agent 1's signal s_1 and his action b_1 , instead she conditions her choice on her private signal s_2 only. Thus, the maximization problem with naive beliefs is analogous to that of agent 1 and leads to the same optimal bidding:

$$b_2^N = \arg \max_{b_2} E[(\theta - p_2)X_2 | s_2] \quad (10)$$

$$= E[\theta | s_2] = \underline{\theta} + (\bar{\theta} - \underline{\theta})s_2 \quad (11)$$

The naive strategy is depicted as the straight line in Figure 1. Its underlying naive belief can be viewed as stemming from level-k reasoning or from fully cursed beliefs (Eyster and Rabin, 2005). In the level-k framework (for a formulation with private information, see e.g. Crawford and Iriberri (2007)) level-0 players are naturally defined as bidders who ignore their incentives and randomize bids between $\underline{\theta}$ and $\bar{\theta}$. If so, an agent 2 with naive beliefs would be equivalent to a level-1 agent who assumes playing against a level-0 opponent. In the current setting, this also coincides with a 'fully cursed' strategy (Eyster and Rabin, 2005): fully cursed agents believe that other agents play a constant average strategy regardless their information.⁷

The cursed equilibrium model allows for intermediary levels of cursedness: subjects may not fully, but only partially ignore the information revealed by other agents' actions. Our analysis of Subsection 4.3 is analogous to this in that it allows for milder versions of information neglect.

3.3 Hypotheses

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

Hypothesis 1 *Naive bidding is more prevalent in treatment PL than in treatment S.*

In our second hypothesis, we compare the two simultaneous treatments. Given that both the naive best response and the empirical best response are monotonic, we expect most subjects to follow a monotonic strategy. Modifying the action space from a continuous strategy to a discrete (and therefore possibly non-monotonic) order schedule should not significantly affect the results.

⁷More precisely, assume that agent 2 disregards the possibility that agent 1 bids a dominated value, i.e. she believes that he never bids outside of $[\underline{\theta}, \bar{\theta}]$, and otherwise her belief is fully cursed: she believes that agent 1 bids according to (7) such that the mixture of bids corresponds to bidding the prior expectation $\bar{b}_1 = \int b^*(s_1)f(s_1) ds_1 = E[\theta]$.




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

Hypothesis 2 *On average, bidding behavior is identical between treatments PL and LO.*

The final hypothesis considers the possibility that *all* participants acting as agent 2 have naive beliefs. In this case, we would not observe any systematic difference between agent 2's bids and those of agent 1 who observes private information only. We can therefore use the bidding behavior of agent 1 as an empirical benchmark for naive agents 2. For this comparison, we restrict to treatment LO, where both agents have identical action sets.

Hypothesis 3 *In treatment LO, bids of participants acting as agent 2 do not significantly differ from bids of participants acting as agent 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 186 students are recruited with the laboratory's ORSEE database (Greiner, 2004). 62 participants are in each of the treatments PL, S and LO, each with three sessions of 20 or 22 participants. Within each session, the participants are divided into two equally sized groups, the pools of agents 1 and agents 2. Participants remain in the same group throughout the session and repeat the game for 20 periods. At the beginning of

each period, participants of group 1 and 2 are uniform-randomly matched into pairs and the game commences with Nature's draw of θ , followed by the remainder of each game as described in Section 2. At the end of each period, participants 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.

For a complete understanding of the experiment, participants read the instructions for both groups before learning which group they are assigned to. The instructions include a computer-based simulation of the signal structure as well as an understanding test.⁸

The games' parameters have the following values in the experiment, where units of experimental currency are converted to money by a factor of EUR 0.07 per unit:⁹

$$\begin{array}{ll} \underline{\theta} = 40 & \bar{\theta} = 220 \\ \underline{p} = 30 & \bar{p} = 280 \\ \gamma = 0.7 & \end{array}$$

Each session lasted approximately 90 minutes and students earned on average 25.56 Euros. Total earnings consist of a show-up fee of 5 Euros, an endowment of EUR 17.00 and profits from the randomly drawn period (which could be negative but cannot deplete the entire endowment).


4.2 Descriptive Statistics

4.2.1 Behavior of Agent 1: Mild Over-Reaction to Signal

Figure 2 shows the limit orders of participants acting as agent 1, with their corresponding private signal on the horizontal axis. The figure also includes the theoretical prediction (thick black line) and the results of quartile regressions. The figure shows that the median bids increase slightly but significantly faster in the signal than is predicted by the benchmark theory: agent 1 bids, on average, more aggressively than predicted. For our analysis, this implies that the empirical best response (EBR) of agent 2 differs from the theoretical Bayes Nash best response of agent 2: EBR is closer to the naive prediction. This leaves somewhat less room for identifying a deviation of agent 2 in the direction of naive bids. But as the next subsections show, such identification is still possible.

⁸Instructions can be provided upon request.

⁹The range for the limit order and the price [30; 280] is chosen asymmetrically around 40 and 220 in order to make the strategy of always choosing the mean $130 (= \frac{1}{2} \cdot (40 + 220))$ less salient.



Notes: Regression lines depict a quantile regression with robust standard errors clustered by participant. Regression is computed pooling observations of all treatments.

Figure 2: Quantile regression of bids in group 1.

4.2.2 Behavior of Agent 2: Differences between Treatments PL and S.

This subsection examines descriptive statistics regarding Hypothesis 1. The hypothesis implies a tendency of low prices to induce more sell decisions in treatment S than treatment PL, and of high prices to induce more buys. Altogether, the hypothesis predicts a more linear relationship between signals and reservation prices in treatment PL than in treatment S. For data summaries that are relevant for Hypotheses 2 and 3, we refer to the following subsections.

Figures 3 and 4 depict transaction prices and signals in treatments PL and S, respectively. For each price-signal observation, a black marker indicates a buy, whereas a white marker indicates a sell. In treatment PL, acquisitions and sales separate clearly between the upper-left and lower-right portions of the figure. In contrast, the data from treatment S show a more blended set of decisions. As predicted by Hypothesis 1, treatment S induces more buys at high prices and more sales at low prices. The separation between black versus white markers is more blurred but appears to be steeper, as predicted by Hypothesis 1. At least on average, participants account more for the realized price in treatment S.

We now quantify the effect and also 'zoom in' on interesting cases: we consider whether information contained in the price is especially important in situations where the participants' own signal is relatively uninformative. In these cases, Hypothesis 1 should be especially relevant. Table 2 reports the shares of




Figure 3: Sells vs. buys in treatment PL.




Figure 4: Sells vs. buys in treatment S.

buys and sells for all observations, and separately for observations where the signal lies in $[0.4, 0.6]$.¹⁰ To account for the symmetry of buys and sells, the table separates the data according to whether the price lies below (left panel) or above (right panel) its theoretical expectation of 137.5.¹¹

Table 2: Share of buys at high and low prices^a

	$p_2 \leq 137.5$		$p_2 > 137.5$	
	All signals	[0.4 - 0.6]	All signals	[0.4 - 0.6]
Treatment S	.4477 (.038)	.5072 (.063)	.6341 (.038)	.6102 (.076)
Treatment PL	.5029 (.031)	.6957 (.069)	.35 (.029)	.0847 (.038)
Difference	-.0552	-.1885*	.2841***	.5255***

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Participant-clustered standard errors in parentheses.

^a Shares of buys are computed with respect to the market price of treatment S: the share in treatment PL represents the participants' hypothetical choices that would come about if the set of market prices of treatment S (rounded to the nearest ten) were also realized in treatment PL. See table A3 and A4 in appendix A for a comparison with actual treatment prices.

As the table shows, observing a low price induces participants to sell only mildly more in treatment S than in treatment PL (55% versus 50%, see the first column). The treatment effect is much stronger, however, when prices are high: participants of treatment S show a much stronger tendency to buy if $p_2 > 137.5$. The choice ratios even reverse between the treatments: 0.37/0.63 in treatment S versus 0.65/0.35 in treatment PL. As expected, the difference between both treatments increases when narrowing the focus to relatively uninformative signals between 0.4 and 0.6. In these cases, participants buy at high prices with far higher frequency in treatment S than in treatment PL, 0.61 compared to 0.08.¹²

Table 3 narrows the focus to cases where the price contains information that goes against the agent's signal. In these cases, the treatment difference is also highly significant. In treatment S, we observe a share of 0.29 buys, among all decisions with a high price but with an own signal below 0.5. In treatment PL this share is much lower, at a mere 0.03. Once again, this pattern of acting against

¹⁰Tables A3 and A4 in Appendix A contain a data split with a larger number of categories.

¹¹ $E[p_2] = 0.7E[b_1] + 0.3E[\varepsilon] = 137.5$ holds if agent 1 chooses his risk-neutral strategy b_1^* .

¹²Median earnings in treatment S are about 10 percent higher than in the simultaneous treatments (see table A6 in the Appendix).

Table 3: Acting against one's own signal^a

	$p_2 \leq 137.5$ $s_2 > .5$	$p_2 > 137.5$ $s_2 \leq .5$
TMT S (Seq.)	.8043 (.041)	.2870 (.060)
TMT PL (Sim.)	.9118 (.024)	.0364 (.017)
Difference	-.1075**	.2506***

Note: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$. Participant-clustered standard errors in parentheses.

^a Share of buys is computed with respect to the market price of the treatment S (see note of Table 2).

one's own signal suggests that participants learn from the price only in treatment S. In the table's left column, we observe the analogous pattern for low prices and high own signals: in treatment S there is a substantially larger proportion of participants following the price (0.20 versus 0.09).

4.3 Testing Hypotheses

Hypothesis 1. To evaluate the degree of 'cursedness', we focus on those areas of Figures 3 and 4 where naive and optimal strategies make different predictions. Within this area, we measure the proportion η of naive decisions:

$$\eta = \frac{b_N}{b_N + b_B} \quad (12)$$

where b_N and b_B denote the number of orders consistent with naive and EBR predictions, respectively.

Figures 5 and 6 show the relevant observations in treatment PL and treatment S, respectively. For these observations, naive expectations induce buys for signals below 0.5 and sells for signals above 0.5, while Bayesian expectations induce opposite actions. The empirical measures b_N and b_B correspond to the number of triangle markers and cross markers, respectively. Hypothesis 1 corresponds to the prediction that the proportion of naive choices is larger in treatment PL than in treatment S: $\eta^{PL} > \eta^S$.

Table 4 shows that the simultaneous treatments LO and PL show equally high proportions of naive decisions, at 0.42 of the relevant observations. In treatment S, this share is substantially reduced to 0.16. The difference between the PL and S treatments is significant ($p < 0.001$, Wald test).




Figure 5: Sells and buys within the relevant area in treatment PL.




Figure 6: Sells and buys within the relevant area in treatment S.

In Appendix A, we also consider the evolution of decisions in the course of the experiment. Although behavior is relatively stable over the 20 repetitions, we observe that learning success is better in treatment S than in treatments PL and LO. If anything, the trading decisions become less rational in the two simultaneous treatments.

Table 4: Proportion of naive decisions

	LO	PL	S
η	.4224 (.0606)	.4231 (.0691)	.1553 (.0389)
N	116	104	103

Note: CRSE in parentheses.


Table 5: Share of buys for varying price/signal areas

Treatment	$p_2 \leq 137.5$		$p_2 > 137.5$	
	$s_2 \leq 0.5$	$s_2 > 0.5$	$s_2 \leq 0.5$	$s_2 > 0.5$
LO	.2686 (.037)	.9603 (.023)	.0344 (.015)	.5454 (.052)
PL	.2112 (.040)	.8705 (.030)	.0384 (.018)	.5609 (.051)
Difference	.0574	.0898**	-.004	-.0155

Note: *p < 0.1, **p < 0.05, ***p < 0.01.

Hypothesis 2. The hypothesis compares the simultaneous treatments LO and PL. Their equal proportions of naive decisions (Table 4) already suggests that bidding is very similar in the two treatments. A direct comparison between participants is, however, hampered by the fact that the strategy spaces differ between the two treatments. For a primary comparison, we neglect the different strategy spaces and compare the overall buy/sell frequencies. As a second step, we examine whether the two treatments induce differences in monotonic strategies.

To compare the overall buy/sell frequencies behavior, we report in Table 5 the shares of buys for four different areas of prices and signals. Shares of buys are roughly the same in all areas, and the differences are insignificant in all but one area: for low prices and high signals the amount of buys differs by 9 percentage



Polynomial random effects regression.


Figure 7: Expected bids in treatments PL and LO.

points. However, this difference is not driven by non-monotonic strategies of participants in treatment PL. In 89% of cases, participants in treatment PL exhibit a unique switching point that separates their buys and sells. These strategies correspond to monotonic cutoff strategies and we can directly compare them with the data from treatment LO.¹³ When confining the analysis to monotonic strategies, a significant difference of 7 percentage points remains in the relevant price-signal region. Yet, this difference is not reflected in the average bidding functions. Figure 7 shows the result of two 3rd-order polynomial regressions, one for each of treatments PL and LO. The estimated bidding curves are very close, except for extreme values of the signal, near 0 or 1. The differences are mainly driven by the larger amount of dominated choices observed in treatment LO, where 20.5% of all limit orders lie strictly below 40 or strictly above 220 (compared to 0.6% of the monotonic strategies in treatment PL). The two bidding curves do not statistically differ (see Table A1 in Appendix A for a more detailed comparison).

Hypothesis 3. Hypothesis 3 compares the limit orders of agent 1 and agent 2 in treatment LO, where the actions sets are identical between the two agents.

The two polynomial regressions reported in Figure 8 reveal that average bid functions of both agents do not differ significantly. For details of the comparison,

¹³A further minor difference is that the inferred cutoff values in treatment PL lie on a grid, where the limit orders in treatment LO were unrestricted within $[p, \bar{p}]$. We interpolate the responses from treatment PL: e.g. if a participant buys up to $p_2 = 70$ and sells for $p_2 \geq 80$, we take 75 as the cutoff.



Polynomial random effects regression.

Figure 8: Bids in groups 1 and 2 in treatment LO.

see Table A2 in appendix A.

Summarizing the evidence on Hypotheses 2 and 3, we find that in the simultaneous treatments the participants acting in the role of agent 2 behave in a way that cannot be distinguished from participants acting in the role of agent 1: despite the fact that the price carries information about the fundamental value, they behave as if they only take account of their private signals. The following subsection further examines this finding, by quantifying the reaction to each piece of information via a random utility model.

4.4 Random Utility Model

We assume the participant-specific buying probability of agent 2 to follow a logistic distribution but allow for an over-weighted or under-weighted relevance of the available pieces of information:

$$P(X_i = +1|u_i, s_2, p_2) = \frac{e^{\lambda(\hat{E}[\theta|s_2, p_2] - p_2 + u_i)}}{1 + e^{\lambda(\hat{E}[\theta|s_2, p_2] - p_2 + u_i)}} \quad (13)$$

with

$$\hat{E}[\theta|s_2, p_2] = 40 + 180 \cdot \hat{P}(\theta = 220|s_2, p_2) \quad (14)$$

$$\hat{P}(\theta = 220|s_2, p_2) = [1 + LR(s_2)^{-\beta} \cdot LR(p_2)^{-\alpha}]^{-1} \quad (15)$$

The choice probability (13) depends on subjective expected payoff, which in our case is the difference between a subjective belief $\hat{E}[\theta|s_2, p_2]$ and the market price p_2 . λ is the logistic precision parameter and u_i is participant i 's random utility shifter, which we assume as i.i.d. normally distributed with mean 0 and variance σ_u^2 . To allow for irrational weighting of information, $\hat{P}(\theta = 220|s_2, p_2)$ represents the subjective posterior probability of the event that $\theta = 220$, given (s_2, p_2) . This posterior probability depends on the likelihood ratio of the signal, $LR(s_2) \equiv \frac{P(\theta=220|s_2)}{P(\theta=40|s_2)}$, and on the likelihood ratio of the price $LR(p_2) \equiv \frac{P(\theta=220|p_2)}{P(\theta=40|p_2)}$. β and α are the potentially irrational weights that the participant assigns to the signal's and the price's informational content, respectively. 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 lying between zero and one. A rational trader, on the other hand, 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, captured by the case that either β or α are greater than one.

Table 6: Results of MSL estimation

	Treatment PL	Treatment S
β	1.927 (0.96)	1.374 (0.55)
α	0.441* (0.43)	1.879** (0.41)
λ	0.0391** (0.012)	0.0393*** (0.011)
σ_u	0.011	0.048
N	620	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.

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 $P(p_2|\theta = 40)$ 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.

The estimates are reported in Table 6 and confirm the findings of the previous subsections. In treatment PL, the estimated α -coefficient of 0.44 lies sig-

nificantly below the optimal value 1, suggesting that participants in group 2 pay too little attention to the price's informativeness. Instead, they over-weight their private signals ($\hat{\beta} = 1.93$). In treatment S, the perceived levels of informativeness of signal relative to price are reversed: this treatment induces a significant over-weighting of the price's likelihood ratio ($\hat{\alpha} = 1.87$).¹⁴ The importance of the private signal decreases from 1.93 to 1.37 between treatments PL and S. Both β estimates do not significantly differ from 1.

5 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 detect covert information while focusing on a single outcome. The degree of inference and consequently the quality of informational efficiency interact with market design.

¹⁴This 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.

References

- 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. ISSN 0022-0027.
- J. D. Carrillo and T. R. Palfrey. No trade. *Games and Economic Behavior*, 71(1): 66--87, Jan. 2011. ISSN 08998256.
- 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.
- 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.
- I. Esponda and E. Vespa. Hypothetical Thinking and Information Extraction in the Laboratory. *American Economic Journal: Microeconomics*, 6(4):180--202, 2014.
- J. S. B. T. Evans. *Hypothetical Thinking: Dual Processes in Reasoning and Judgment*. 2007. ISBN 1135419531.
- E. Eyster and M. Rabin. Cursed equilibrium. *Econometrica*, 73(5):1623--1672, 2005.
- U. Fischbacher. Z-Tree: Zurich toolbox for ready-made economic experiments. *Experimental Economics*, 10(2):171--178, 2007. ISSN 13864157.
- 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.
- 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 1465-7295.
- 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.
- 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.
- 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. ISSN 00305073.
- 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.
- E. Shafir and A. Tversky. Thinking through uncertainty: nonconsequential reasoning and choice. *Cognitive psychology*, 24(4):449--474, 1992. ISSN 00100285.

A Appendix

A.1 Regression Results

Table A1: Pooled regression with group 2 of treatments LO and PL.

	Coeff.	SE
s_i	-156.4**	(53.75)
s_i^2	846.3***	(145.3)
s_i^3	-526.8***	(98.76)
$s_i \cdot LO$	17.45	(84.84)
$s_i^2 \cdot LO$	-26.75	(229.4)
$s_i^3 \cdot LO$	25.13	(152.9)
LO	-4.037	(5.493)
_cons	50.55***	(3.569)
N	1169	
σ_u	13.57	
σ_ϵ	32.22	
ρ	0.151	

Note: Robust standard errors statistics in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The dummy variable LO takes the value 1 for treatment LO. The Wald test of the joint hypothesis that LO and all its interaction terms have zero coefficients results in a p-value of 0.5087.

Table A2: Pooled regression for treatment LO.

	Coeff.	SE
s_i	-79.04	(52.19)
s_i^2	688.1***	(137.5)
s_i^3	-418.2***	(92.20)
$s_i * A2$	-59.91	(83.86)
$s_i^2 * A2$	131.5	(224.5)
$s_i^3 * A2$	-83.51	(148.7)
$A2$	5.275	(5.522)
$_cons$	41.22***	(3.620)
N	1240	
σ_u	13.48	
σ_ϵ	32.56	
ρ	0.146	

Note: Robust standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

A2 is a dummy variable for agent 2. The Wald test of the joint hypothesis that A2 and all its interaction terms have zero coefficients results in a p-value of 0.8517.

A.2 Descriptive Statistics

Table A3: 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]
S (Seq.)	.4477 (.038)	.1613 (.061)	.1549 (.050)	.5072 (.063)	.8 (.057)	.8824 (.036)
PL (Sim.)	.5029 (.031)	.1196 (.035)	.1571 (.042)	.6957 (.069)	.9492 (.029)	.9 (.047)
PL (Sim.) (treatment prices)	.4716 (.035)	.1042 (.042)	.1528 (.039)	.6087 (.076)	.8676 (.037)	.9362 (.034)
Hypothetical difference	-.0552	.0417	-.0022	-.1885*	-.1492*	-.0176
Actual difference	-.0239	.0571	.0021	-.1015	-.0676	-.0538

Note: CRSE in parentheses. Share of buys is computed with respect to the market price of the sequential treatment, unless treatment prices is denoted in parentheses. In the latter case, we use the transaction prices that occurred in the specific treatment. The actual difference refers to the difference between the shares computed with actual transaction prices, whereas hypothetical difference is obtained with prices of treatment S for both treatment shares.

Table A4: 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]
S (Seq.)	.6341 (.038)	.0333 (.032)	.2391 (.069)	.6102 (.076)	.8919 (.037)	.9104 (.049)
PL (Sim.)	.35 (.029)	0 (.)	.0213 (.021)	.0847 (.038)	.4933 (.060)	.8088 (.046)
PL (Sim.) (treatment prices)	.3582 (.039)	0 (.)	.0222 (.021)	.1186 (.048)	.5455 (.075)	.7324 (.064)
Hypothetical difference	.2841***	.0333	.2178**	.5255***	.3986***	.1016
Actual difference	.2759***	.0333	.2169**	.4916***	.3464***	.178*

Note: CRSE in parentheses. Share of buys is computed with respect to the market price of the sequential treatment, unless treatment prices is denoted in parentheses. In the latter case, we use the transaction prices that occurred in the specific treatment. The actual difference refers to the difference between the shares computed with actual transaction prices, whereas hypothetical difference is obtained with prices of treatment S for both treatment shares.

Table A5: Acting against one's own signal (treatment prices)^a

	$p_2 \leq 137.5$ $s_2 > .5$	$p_2 > 137.5$ $s_2 \leq .5$
TMT S (Seq.)	.8043 (.041)	.2870 (.060)
TMT PL (Sim.)	.8705 (.030)	.0392 (.019)
Difference	-.0662	.2478***

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

^a Share of buys is computed with respect to the transaction prices in the relevant treatment.

Table A6: Profits of participants in group 2

	TMT LO (Sim.)	TMT PL (Sim.)	TMT S (Seq.)
Mean	33.62	31.89	29.67
Median	34.85	35	39.65
Standard dev.	72.84	75.15	76.12

Note: The mean in treatment S is considerably lower than the median due to some outliers: a steeper bidding curve in treatment S entails small probabilities of severe losses.

A.3 Individual Behavior

Similar to Carrillo and Palfrey (2011), we classify participants into different categories to ensure that the results are not mainly driven by few participants. We define participants as Bayesian or naive if the number of decisions consistent with the one or the other prediction surpasses a specific threshold. For a participant to be classified, the threshold criterion must apply to the number of decisions both in- and outside the area where the two predictions differ (henceforth, the relevant area). To those who are classified neither as Bayesian nor as naive belong, on the one hand, those who do not clearly favor naive over Bayesian decisions and vice versa, but also those whose decisions are often not consistent neither with the best response nor with the naive model. For instance, participants who overweight the price's relatively to the signal's informativeness may sell at low prices and high signals, although in that case both models predict a buy. We consider three different thresholds: 65%, 75% and 80%.

Table A7: Classification of participants in group 2

Threshold		LO (Sim.)	PL (Sim.)	S (Seq.)
65%	naive	11	10	1
	Bayesian	12	19	26
	neither...nor...	8	2	4
75%	naive	8	5	1
	Bayesian	10	14	19
	neither...nor...	13	12	11
80%	naive	6	5	1
	Bayesian	9	11	17
	neither...nor...	16	15	13
	N	31	31	31

All treatments display more Bayesian than naive participants, but the relative majority of Bayesian participants increases when comparing LO with PL, and when comparing PL with S (see Table A7). Thus, the higher share of Bayesian decisions in treatment S is mirrored by an increased fraction of Bayesian subjects in this treatment.

A.4 Learning

To investigate whether participants learn over time, we divide the observations into two time subsections: an early time interval for the rounds one to ten and a late interval for later rounds.

Participants in the simultaneous treatments LO and PL do not show differences in behavior over time. However, participants of treatment S improve their average bidding behavior towards the empirical best response in the last ten rounds. This improvement is significant if we include all decisions in the analysis. The corresponding Wald test rejects identical parameters in the 3rd-degree polynomial random effects regression with a p-value of 0.0747. The corresponding regression results are depicted in Figure 9. In this sense, the sequential variant of the game enables the participants to learn 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 any of the individual treatments, as shown in Table A8. However, the table shows that in the sequential treatment, participants decide more rationally over time whereas they decide less rationally in the other two treatments. This difference in differences is significant: bootstrapping a distribution under zero difference in differences yields a p-value of 0.01 for the divergence of -20.22 percentage points between treatments PL and S. Analogously, the difference between treatments LO and S is significant with a bootstrap p-value of 0.06.

Table A8: Proportion of naive decisions

	LO	PL	S
First 10	.3774 (.073)	.3654 (.074)	.2 (.071)
Last 10	.4603 (.074)	.4808 (.087)	.1132 (.046)
Difference	-.0830	-.1154	.0868
N	116	104	103

Note: CRSE in parentheses.




Figure 9: Polynomial regression for early and late rounds for the treatments LO, PL and S, respectively.