DISSERTATION

Decision-Making in Markets

zur Erlangung des akademischen Grades
doctor rerum politicarum
(Doktor der Wirtschaftswissenschaft)

eingereicht an der

Wirtschaftswissenschaftlichen Fakultät
der Humboldt-Universität zu Berlin
von

Dipl.-Volksw. Kathleen M. Ngangoue

Präsidentin der Humboldt-Universität zu Berlin:
Prof. Dr.-Ing. Dr. S. Kunst

Dekan der Wirtschaftswissenschaftlichen Fakultät:
Prof. Dr. C. Schade

Gutachter: 1. Prof. G. Weizsäcker, Ph.D.
2. Prof. A. Guarino, Ph.D.

Tag des Kolloquiums: 17. Juli 2017

Decision-Making in Markets

Kathleen M. Ngangoué

June 2017

Pour Bel-Ami.

Acknowledgments

I was blessed to have Georg Weizsäcker as my advisor. More than just an advisor, over the years he became an inspiring mentor, whose opinion and sincere criticism I highly appreciate. With patience, sensitivity, and many discussions, he introduced me to research, which I would not have enjoyed as much without his constant support.

I had the pleasure to meet my second advisor Antonio Guarino during my research visit at University College London. He supported me with much of his time and great care, teaching me so much that I decided to extend my visit. Ever since, he has been a precious advisor on various matters.

I thank Pio Baake, Yves Breitmoser, and Steffen Huck, a trio of experts I often approached to discuss my research. All of my work benefited from their insights and critiques. In addition, I have learned from discussions with Dirk Engelmann, Frank Heinemann, Sveta Ivanova-Stenzel, and Johanna Möllerström. In particular, they have all been involved in developing the Berlin Behavioral Economics network, which provides frequent, interesting opportunities for exchanging ideas.

I gratefully acknowledge financial support from the European Research Council (Starting Grant 263412), the German Science Foundation, Humdoldt University Berlin, and DIW Berlin. I also thank colleagues at Berlin Social Science Center and Humboldt University Berlin for the support in conducting the experiments. Furthermore, the DIW Graduate Center provided, with its masterclasses, travel funding, and seminars, a unique infrastructure for the education of young economists. The flourishing of this young program is mainly due to the Deans Georg Meran, Helmut Lütkepohl and Georg Weizsäcker, as well as the whole Graduate Center team.

The last years would not have been memorable without my colleagues who I am happy to now call my friends, in particular Alexandra Avdeenko, Ludwig Ensthaler, Holger Lüthen, Stephanie Schneider, Doreen Triebe, Paul Viefers, and my dear officemate Christian Zankiewicz, who endured my busy days with a lot of humor.

Finally, I am indebted to my closest friends and to my family, who gave me the strength and freedom to persevere in my research–especially my fantastically

supportive spouse Attila, with whom I shared every single day of this itinerary.
Un très grand merci à vous, mes chers, vous êtes formidables.

Summary

This dissertation studies various channels through which information processing affects market decisions.

Markets have the essential function of aggregating information that is dispersed among participants. In this respect, markets that fully aggregate information are efficient in that they enable participants to easily access necessary information.

A considerable theoretical and empirical literature examines the extent to which markets aggregate information.[1] All in all, informational efficiency is sensitive to technicalities of market designs, including information costs (Grossman and Stiglitz, 1980), transparency (i.e. common knowledge of preferences in Forsythe and Lundholm 1990), and pricing rules (Pouget, 2007). While mis-aggregation can be partly rationalized by frictions (e.g. due to limited arbitrage), it is not always compatible with standard economic theory. Standard economic theory assumes rational decision-makers, who perfectly process all available information. It does not account for trading decisions varying with market features that are irrelevant to optimal decision-making.

If anything, information aggregation relates to participants' ability to process information. For instance, in the asset market experiments of Plott and Sunder (1988) and Forsythe and Lundholm (1990), trading experience improves price discovery.

The extent to which agents accurately process information depends on exogenous factors like the nature of information or the decision context (Olshavsky,

[1]Important early contributions are, among others, Fama (1970); Grossman (1976, 1978); Radner (1979); Grossman and Stiglitz (1980); Hellwig (1980); Diamond and Verrecchia (1981); Glosten and Milgrom (1985); Plott and Sunder (1988); Forsythe and Lundholm (1990).

1979; Gigerenzer and Hoffrage, 1995; Evans, 2007). In addition, decision-makers' inherent abilities generate differences in information usage. This dissertation analyzes these two aspects of information processing in markets. First, it seeks to identify decision contexts that systematically affect belief revision and, consequently, market behavior. Second, it explores whether human inference correlates with inherent risk attitudes.

Analyzing decision-making in markets is, however, a delicate endeavor. Various factors flow into market decisions. Beliefs, preferences, and strategic considerations in the interaction with other participants determine traders' best responses. To control for non-relevant determinants, trading decisions are studied in individual decision-making experiments. Moreover, as market functioning is the main interest, the focus in all experiments lie in trading decisions as opposed to belief elicitation. To this end, the experiments study under different conditions how decisions vary with belief revision. Comparing decisions across treatments enables then to deduce the effect of information processing.

The first chapter, based on joint work with Georg Weizsäcker, investigates how traders process information contained in realized versus hypothetical stimuli. We focus on markets with diverse information and informative prices. More precisely, agents receive private signals and trade at prices that reveal additional information held by other market participants. Agents with rational expectations condition their beliefs on both their private information and the information contained in the price. However, conditioning on observable prices requires a different level of sophistication than conditioning on future prices, which demands more forward-thinking reflection.

In view of this, it is noticeable that market mechanisms differ with respect to whether they require hypothetical thinking. In simultaneous markets, for instance, subjects submit their trading strategies before knowing prices or other participants' actions. Thus, trading strategies are defined for all contingencies. In sequential markets, on the other hand, they know prices at which trade might occur and, therefore, specify their preferences for a single price. Hence, different trading mechanisms demand different levels of cognitive sophistication. This difference is irrelevant under rational expectations, but possibly accounts for deviations from optimal decisions observed in simultaneous markets, like in sealed-bid auctions (Levin et al., 1996).

We explore the relevance of hypothetical thinking in information processing by comparing investment decisions in simultaneous and sequential markets. Under rational expectations, the two market mechanisms in our experiment are isomorphic with respect to strategies and payoffs and should therefore entail the same investment decisions. The results, however, reveal that hypothetical thinking impedes subjects' ability to consider implicit information. In simultaneous markets, where subjects submit their trading strategies before knowing the stimulus (here the price), subjects tend to neglect information contained in prices. In sequential trading mechanisms, where subjects observe stimuli first, they consider price's informativeness and make more rational decisions.

This information neglect is robust to various treatment variations. For instance, removing strategic uncertainty by setting a pre-defined pricing algorithm does not eliminate the bias. Making information in the price more salient and more important by increasing its precision relative to private information also does not reduce the bias. In addition, we explore the underlying mechanisms that make learning from hypothetical events so difficult. The difficulties in hypothetical thinking may stem from two sources. First, anticipating information from an event that takes place in the future requires sophisticated forward thinking. Second, hypothetical thinking entails reflecting about several contingencies. That is, the number of possible outcomes dictates the computational, cognitive challenge. A treatment variation simplifies the decision context by reducing the dimension of possible prices, but retains the difficulty of conditioning on a future price. Results in this treatment disclose that difficulties in hypothetical thinking cannot be assigned to a single source. Reducing the dimension of hypothetical outcomes improves trading decisions, but does not fully eliminate information neglect in simultaneous markets.

In sum, Chapter one shows that the nature of information matters for information usage. Yet, it also shows that market behavior can be improved with a mechanism design, in which implicit information is revealed by realized rather than hypothetical stimuli. Market design, however, may have limited effects on other determinants of information processing, like the type of uncertainty that makes information valuable. The second chapter compares information processing for decisions under risk and ambiguity. While under risk, investors know the distribution of states of nature, ambiguity describes the lack of such precise knowledge.

Despite ambiguity being the more natural setting of uncertainty, fundamental concepts like Rational Expectation Equilibria are defined for risky markets only. One of the reasons is that little is known about how subjects learn under ambiguity. Even with rational expectations, there is no unique benchmark. Theoretically, belief revision under ambiguity could result in either extreme sensitiveness to information, or a conservative update. The utilized updating rule, in turn, determines the speed at which prices converge to true values. Chapter two discloses how belief revision affects decisions in ambiguous markets.

The experiment is implemented with two treatments in a two-by-two design, where decisions vary along the dimension of uncertainty and along the dimension of belief revision. In one treatment, subjects make investment decisions under both risk and ambiguity, but do not need to revise beliefs. In the second treatment, subjects make decisions under risk and ambiguity as well, but have to revise beliefs after the arrival of an informative signal. If information processing is similar under ambiguity and risk, belief revision should not alter the general difference between investments under risk and ambiguity.

As expected, decisions in risky and ambiguous markets differ: Subjects reduce their market participation in ambiguous markets relative to risky markets, and display thereby ambiguity aversion. Learning, on the other hand, does not affect this difference: Average market participation is the same with and without belief revision. Learning alters, however, the quotes at which investors are willing to trade: subjects tender more extreme bids and asks. The results feature heterogeneity in decisions, but are, on average, consistent with a model of recursive smooth preferences. A consequence of recursive preferences is that breaking down information into pieces yields more extreme average beliefs. Hence, with rational recursive preferences, information processing under ambiguity differs from learning under risk in the sense that the frequency of information release determines how quickly prices converge toward true values.

On the other hand, the speed of price discovery also depends on heterogeneity in trading decisions. The degree of heterogeneity might even alter equilibria in some markets (Chapman and Polkovnichenko, 2009; Haltiwanger and Waldman, 1985, 1989). In the experiment, trading decisions are heterogeneous, even when controlling for information. In such markets where agents are symmetrically informed, heterogeneity emanates from random errors, differences in preferences, or differences in belief formation. To understand the origins of heterogeneity, we need to gauge the importance of and the correlation between single determinants. Over and above interpreting aggregate statistics, Chapter three estimates the correlation between preferences and decisions. The analysis combines data on investment decisions under risk and data on elicited risk preferences. Even after correcting for measurement error in elicited preferences, risk preferences have a moderate weight in investment decisions. The weight is, however, amplified when learning occurs. As the correlation between preferences and decisions varies across the information condition, the interaction between learning abilities and risk preferences is also studied. Testing for Bayesian inference reveals that risk-averse subjects update more conservatively, and have therefore a different risk perception in the investment task. Hence, the analysis discloses that variance in preferences determines variance in decisions not only directly, but also through belief revision.

In sum, the individual decision-making experiments in this dissertation point to different aspects of information processing that systematically affect market decisions. The necessity to engage in hypothetical thinking impedes optimal decision-making. A market mechanism that dispenses with hypothetical thinking, however, yields decisions close to the rational benchmark. The environment in which information is obtained matters as well: Under ambiguity, beliefs become more extreme with gradual information release. Over and beyond exogenous factors, the inherent preferences of decision-makers impact their learning abilities and, in this way, determine heterogeneity in decisions after belief revision. Certainly, the findings' robustness to market feedback and effects on market parameters like price volatility or efficiency remain to be explored in future research.

Contents

List of Figures

List of Tables

1 Learning from unrealized versus realized prices

This chapter is based on joint work with Georg Weizsäcker.

1.1 Introduction

Market prices reflect much information about fundamental values. The extent to which traders are able to utilize this information has important welfare consequences but is difficult to measure as one often lacks control of the traders' restrictions, beliefs and preferences. One possibility to detect a bias in price inference is to modify the informational environment in a way that is irrelevant for rational traders. If trading reacts to a framing variation that is uninformative under rational expectations, the latter assumption is questionable. We focus on an important dimension of variability between markets, the conditionality of price. In *simultaneous* markets, the price realization is unknown to the traders at the time when they make their decisions—examples are financial markets with limit orders or other supply/demand function regimes. Theoretically, traders would incorporate the information of each possible price into their bids, as in the Rational Expectations Equilibrium prediction by Grossman (1976), inter alia. But the price information is hypothetical and traders may find it hard to make the correct inference in hypothetical conditions. A host of evidence on Winner's Curse and other economic decision biases is consistent with this conjecture, as is the psychological evidence on accessibility (Kahneman, 2003) and contingent thinking (Evans, 2007).[1] Simultaneous asset markets with price-taking agents are a relevant point in case for such failures of contingent thinking; one that has not previously been researched, to our knowledge. In contrast, *sequential* markets—e.g. many quote-based markets and sequential auctions—have the traders know the price at which they can complete their

[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).

trades. Here, it may still be nontrivial to learn from the price; but both the psychological research on contingent reasoning and the related economic experiments that include treatment variations where simultaneity is switched on and off (Carrillo and Palfrey 2011, Esponda and Vespa 2014; 2016 and Li 2016) suggest that the task is more accessible in a sequential trading mechanism than in a simultaneous one. Our series of experiments confirms this hypothesis, in a simple and non-strategic market environment where agents act as price takers.[2] In such an environment, the failure to learn from the price is especially noteworthy because the price explicitly reflects the asset value, conditional on the available information. To shed further light on the importance of this failure, we study its potential sources and discuss possible implications in financial markets.

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

Markets that represent hybrid versions of order- and quote-driven mechanisms also exhibit important features of simultaneous trading. For example, equity markets with low liquidity may be cleared throughout the day with periodically conducted call auctions; other markets open or close the day's trading via call auctions. Additionally, an increasing flow of retail orders is internalized (Comerton-Forde et al., 2016). These orders are not executed on public exchanges but are executed internally through dark avenues or routed to different exchanges, making it difficult for retail investors to monitor the market condi-

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

tions prior to trade. Thus, even for continuously traded assets the increasing market fragmentation and the increasing speed of trades force (slow) retail investors to post orders without precise knowledge of transaction prices, requiring contingent thinking.

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

Our participants trade a single, risky, common-value asset. To trade optimally, a participant considers two pieces of information: her private signal and the information conveyed by the asset price. The latter is informative because it is influenced by the trading activity of another market participant who has additional information about the asset value. To manipulate the accessibility of the price information, we perform the experiment in two main treatments, simultaneous (SIM) versus sequential (SEQ). In treatment SIM, participants receive a private signal and submit a limit order. If the market price realizes below the limit, the trader buys one unit of the asset, otherwise she sells one unit.[4] Despite the fact that the price has not yet realized, SIM traders would

[3] 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.

[4] 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 1.2. In each

optimally infer the extent to which a high price indicates a high value and, thus, soften the demand's downward reaction to a higher price, relative to the case that the price is uninformative. The possibility that traders may fail to learn from hypothetical prices is examined by comparing to the treatment with sequential markets, SEQ, where the price is known when traders choose to buy or sell. Conditional thinking is not necessary here but treatments SIM and SEQ are nevertheless equivalent: they have isomorphic strategy sets and isomorphic mappings from strategies to payoffs.

Section 1.2 presents the experimental design in detail and Section 1.3 discusses our behavioral hypotheses. We present three benchmark predictions for comparison with the data: first, full naiveté, where the trader learns nothing from the price; second, the Bayes-Nash prediction, where a trader assumes that previous trades are fully rational and accounts for it; and third, the empirical best response that takes into account the actual distribution of previous trades, which may deviate from optimality. We use the latter as our main benchmark for optimality as it maximizes the traders' expected payments. That is, we ask whether naiveté fits the data better than the empirical best response, separately by treatment.

The data analysis of Section 1.4 shows that the participants' inference of information from the price varies substantially between simultaneous versus sequential markets. In SIM, participants often follow the prediction of the naive model, thus showing ignorance of the information contained in the price. Price matters mainly in its direct influence on the utility from trade—a buyer pays the price, a seller receives it. In contrast, in SEQ, where transaction prices are known beforehand, asset demand is signigficantly more affected by the information contained in the price and the large majority of trades are as predicted by empirical best response. Averaging over all situations where the naive benchmark differs from the empirical best response, the frequency of naive trading decisions is twice as high in SIM relative to SEQ, at 38% versus 19%.

Section 1.5 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

treatment, we restrict the trades to a single unit of supply or demand per trader.

the price. We thus repeat the experiment with two treatments where early traders are much better informed than later traders, rendering learning from the price more important and more salient. We find that the replication only exacerbates the differences between simultaneous and sequential markets, both in terms of behavior and payoff consequences. This evidence makes it implausible that the bias is driven by negligence or the lack of salience of the price's informativeness.

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

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

We then combine the different treatments into an aggregate estimation of information use (Section 1.5.4). The analysis of the combined simultaneous treatments shows that relative to empirical best response, the participants under-weight the information contained in the price to a degree that is statistically significant (at $p = 0.09$ in a one-sided test) and that they strongly over-weight their own signal's importance. In the sequential treatments, they over-weight both price and their own signal. Overall, the estimates indicate that traders far under-weight the prior distribution of the asset's value but that they nevertheless learn too little from the price in simultaneous markets.

Taken together, the experiments provide evidence of an interaction between market microstructure and the efficiency of information usage. In the language introduced by Eyster and Rabin (2005), we find that the degree of 'cursedness of beliefs' is higher when the information contained in the price is less accessible: with price not yet realized, traders behave as if they tend to ignore the connection between other traders' information and the price. Aggregate demand therefore decreases too fast with the price. The economic bearing of the effect is further discussed in Section 1.6. We examine the predictions of Hong and Stein (1999) and Eyster et al. (2015) that markets with naive traders, who cannot learn from the price, generate an inefficient and slow price discovery. Naive traders tend to speculate against the price, pushing it back towards its ex-ante expectation also in cases where their own signals are consistent with the direction of price movement. Their erroneous speculation reduces the extent to which the price reveals the underlying value. Confirming this prediction, we simulate a standard price setting rule with our data and find that price discovery is slower in simultaneous treatments than in sequential treatments. Any (hypothetical) subsequent traders can therefore learn less from the price. But naiveté is detrimental not only to later players: also the observed payoffs of our market participants themselves are lower in SIM than in SEQ, albeit not to a large extent.

While we focus on financial 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 participants prefer decision processes with less cognitive strain. They focus on one model, one alternative or one relevant category when reflecting about possible outcomes and their consequences (Evans, 2007; Murphy and Ross, 1994; Ross and Murphy, 1996). They also process salient and concrete information more easily than abstract information (see e.g. Odean 1998 and the literature discussed there).

Several authors before us have pointed out that a possibility to reduce the complexity of learning is to proceed in a sequential mechanism, like in quote-driven markets.[5] Our experiment suggests a specific manifestation of this effect,

[5] 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

namely that drawing the attention to the realized price may enable the decision maker to interpret more easily the information underlying the price. In the related bilateral bargaining experiment by Carrillo and Palfrey (2011), buyers also trade more rationally in a sequential trading mechanism than in a simultaneous one. They process information more easily and exhibit less non-Nash behavior when facing a take-it-or-leave-it price instead of bidding in a double auction. Auction experiments similarly find that overbidding is substantially reduced in dynamic English auctions compared to sealed-bid auctions (Levin et al., 1996). Other contributions suggest that traders may systematically disregard relevant information that is conveyed by future, not yet realized events: overbidding decreases when finding the optimal solution does not necessitate updating on future events (Charness and Levin, 2009; Koch and Penczynski, 2014).[6] Another related study is the voting experiment of Esponda and Vespa (2014) who find that when the voting rules follow a simultaneous game that requires hypothetical thinking, the majority of participants behave nonstrategically, whereas in the sequential design they are able to extract the relevant information from 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 financial 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 the average retail trader makes too much or too little inference from the price.

1.2 Experimental design

The basic framework is identical across treatments, involving a single risky asset and money. A market consists of two traders, trader 1 and trader 2, who each either buy or sell one unit of the risky asset.[7] The asset is worth $\theta \in \{\underline{\theta}, \overline{\theta}\}$, with

of attention results in a loss of acuity" (p.457).

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

[7]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

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

$$f(s_i|\theta) = \begin{cases} 2(1-s_i) & if \quad \theta = \underline{\theta} \\ 2s_i & if \quad \theta = \overline{\theta} \end{cases} \qquad i \in \{1,2\} \tag{1.1}$$

Conditional on θ, the signals of the two traders are independent.
Each trader i faces a separate transaction price p_i. Trader 1's price p_1 is uniformly distributed in $[\underline{\theta}, \overline{\theta}]$ and is uninformative about the fundamental value θ. Trader 1 observes his private signal s_1 and states his maximum willingness to pay by placing a limit order b_1. If p_1 lies weakly below b_1, he buys one unit of the asset. If p_1 strictly exceeds b_1, he sells one unit.[8] By checking an additional box, trader 1 may convert his limit order into a "reversed" limit order. A reversed limit order entails the opposite actions: the trader buys if p_1 weakly exceeds b_1, otherwise he sells. (Only few participants make use of it; we defer the motivation for allowing reversed limit orders to Section 1.2.2.) Let Z_1 denote the indicator function that takes on value 1 if a limit order is reversed. Trader 1's demand is X_1:

$$X_1 = Y_1(1 - Z_1) - Y_1 Z_1 \tag{1.2}$$

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

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

needs to be adjusted by selling or buying one unit of a given asset.

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

1.2.1 Simultaneous treatment (SIM)

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

Participants are informed that the price p_2 reflects the expectation of an external market maker who observes trader 1's buying or selling decision and who assumes that trader 1 bids rationally upon receipt of his signal s_1. Importantly, to avoid any ambiguity in the description, they learn the pricing rule that maps p_1 and the realized value of X_1, denoted by x_1, into p_2:

$$p_2 = \begin{cases} \frac{\overline{\theta}+p_1}{2}, & \text{if } x_1 = 1 \\ \\ \frac{\underline{\theta}+p_1}{2}, & \text{if } x_1 = -1 \end{cases} \tag{1.3}$$

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

1.2.2 Sequential treatment (SEQ)

In treatment SEQ, trader 2 observes the price p_2 as specified in (1.3) *before* making her decision. The game proceeds sequentially, with trader 1 first choosing his (possibly reversed) limit order b_1. As in treatment SIM, his demand X_1 determines the price for trader 2, p_2. Trader 2 observes the realized value of $\{p_1, p_2, s_2\}$ and chooses between buying and selling at p_2.

It is straightforward to check that treatments SIM and SEQ are strategically equivalent. Treatment SEQ allows for four possible strategies contingent on $p_2 \in \{\frac{\underline{\theta}+p_1}{2}, \frac{\overline{\theta}+p_1}{2}\}$: $\{buy, buy\}$, $\{buy, sell\}$, $\{sell, buy\}$ 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. The two strategy spaces are therefore isomorphic.

1.2.3 Payoffs

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

$$\Pi_i = (\theta - p_i)X_i \tag{1.4}$$

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

1.3 Predictions

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

1.3.1 Rational best response

Trader 1 has only his private signal s_1 to condition his bid upon. His optimal limit order b_1^* is not reversed and maximizes the expected profit conditional on s_1. It is easy to show (using the demand function (1.2)) that b_1^* increases linearly in the signal:

$$b_1^*(s_1) = \arg\max_{b_1}\ E[(\theta - p_1)X_1 | s_1] = E[\theta | s_1] = \underline{\theta} + (\overline{\theta} - \underline{\theta})s_1 \tag{1.5}$$

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

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

the following fixed point:[10]

$$b_2^*(s_2) = E[\theta|s_2, p_2 = b_2^*(s_2)] \tag{1.6}$$

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

The Bayes-Nash (BN) strategy of trader 2, however, simplifies to a step function: p_2 reflects the market maker's expectation (see (1.3)), implying that in equilibrium p_2 would make trader 2 indifferent in the absence of her own signal s_2. The additional information contained in s_2 breaks the tie, such that trader 2 buys for $s_2 \geq E[s_2] = \frac{1}{2}$, and sells otherwise.

However, the BN best response is not the most payoff-relevant 'rational' benchmark. In the experiment, participants in the role of trader 1 deviate from their best response b_1^* and participants acting as trader 2 would optimally adjust to it. Their price p_2 is still informative about θ because it reflects s_1, but p_2 does not generally equal $E[\theta|X_1]$ if X_1 is subject to deviations from b_1^*. For a stronger test of naive beliefs, we therefore consider the empirical best response (EBR) to the participants acting as traders 1. The computation of the empirical best response is computed via a numerical approximation to the fixed point equation (1.6).

The two benchmarks BN and EBR are depicted in Figure 1.1 (for the parameters of the actual experiment that are reported in Section 1.4, and using the empirical behavior described in Section 1.5 for the calculation of EBR), together with the naive prediction that we describe next.[11] 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

[10] For a simple proof of this statement, verify that if b_2^* were to violate (1.6) then there would exist realizations of (p_2, s_2) such that p_2 lies in the vicinity of $E[\theta|s_2, p_2 = b_2^*]$ and profits are forgone.

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

level of p_2, EBR requires trader 2 to buy only if she has additional positive information (large s_2).

1.3.2 Best response to naive beliefs

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

$$b_2^N = \arg\max_{b_2} \; E[(\theta - p_2)X_2|s_2] = E[\theta|s_2] = \underline{\theta} + (\overline{\theta} - \underline{\theta})s_2 \tag{1.7}$$

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

1.3.3 Hypotheses

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

[12] 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 (1.5), 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 (2015) also allows for intermediary levels of cursedness, where agents may only partially ignore the information revealed by other agents' actions. Our estimations in Subsection 1.5.4 also allow for milder versions of information neglect.

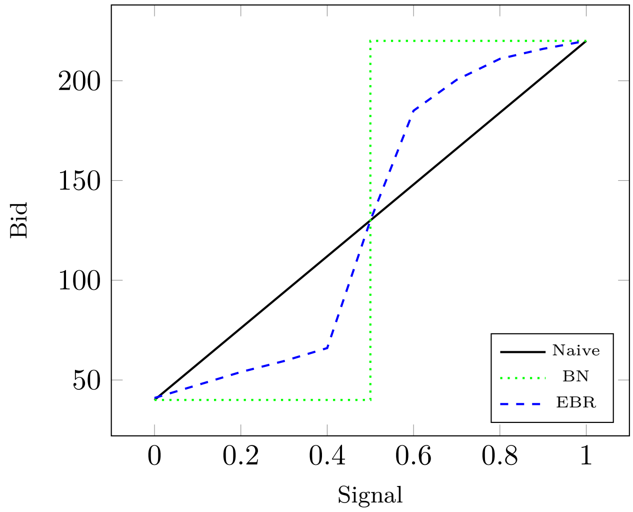

Figure 1.1: Naive, Bayes Nash and empirical best responses.

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 1.1, 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.*

1.4 Experimental procedures and results

1.4.1 Procedures

The computerized experiment is conducted at Technical University Berlin, using the software z-Tree (Fischbacher, 2007). A total of 144 students are recruited with the laboratory's ORSEE database (Greiner, 2004). 72 participants are in each of the treatments SIM and SEQ, each with three sessions of 24 participants. Within each session, the participants are divided into two equally sized groups of traders 1 and traders 2. Participants remain in the same role throughout the session and repeat the market interaction for 20 periods. At the beginning of each period, participants of both player roles are randomly matched into pairs and the interaction commences with Nature's draw of θ, followed by the market rules as described in Section 1.2. At the end of each period, subjects learn the value θ, their own transaction price (if not already known) and their own profit. Upon conclusion of the 20 periods, a uniform random draw determines for every participant one of the 20 periods to be paid out for real.
Participants read the instructions for both roles, traders 1 and 2, before learning which role they are assigned to. The instructions include an elaborate computer-based simulation of the signal structure as well as an understanding test. The support of the asset value is $\{40, 220\}$.[13] Each session lasted approximately 90 minutes and participants earned on average EUR 22.02. Total earnings consist of a show-up fee of EUR 5.00, an endowment of EUR 15.00 and profits from the randomly drawn period (which could be negative but could not deplete the entire endowment). Units of experimental currency are converted to money by a factor of EUR 0.08 per unit.

1.4.2 Results

Trader 1

Figure 1.2 shows the implemented buys and sells of participants acting as trader 1 in treatment SIM, with the corresponding market price on the vertical axis

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

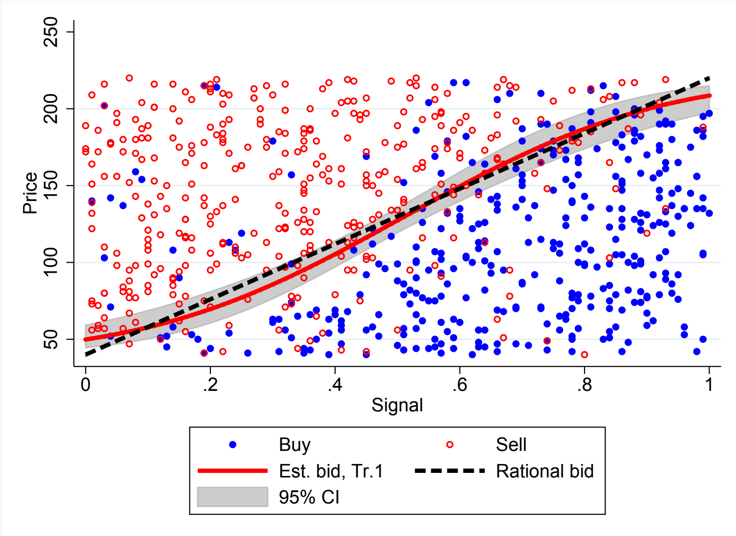

Note: The average bidding curve corresponds to $\underline{\theta} + (\overline{\theta} - \underline{\theta}) \cdot \hat{P}(X_1|s_1)$, where $\hat{P}(X_1|s_1)$ is the probit estimate of the probability of trader 1 buying in treatment SIM.

Figure 1.2: Bids of traders 1.

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.

Trader 2: Testing hypotheses

Hypothesis 1. To evaluate the degree of naiveté, we focus on the area of Figure 1.1 where naive and optimal strategies make different predictions. Within this area, we calculate the proportion η of naive decisions:

$$\eta = \frac{d_N}{d_N + d_B} \tag{1.8}$$

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

Figures 1.3 and 1.4 show the relevant observations in treatments SIM and SEQ, respectively. For these observations, naive expectations induce buys for signals

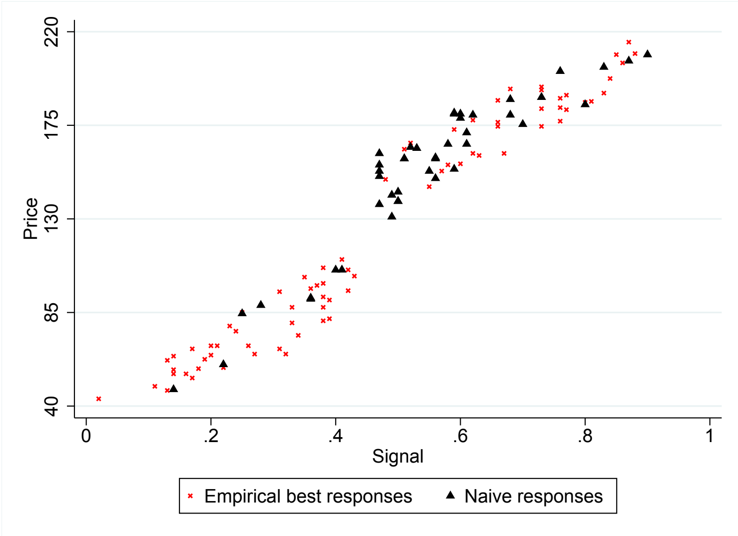

Figure 1.3: Sells and buys within the relevant area in treatment SIM.

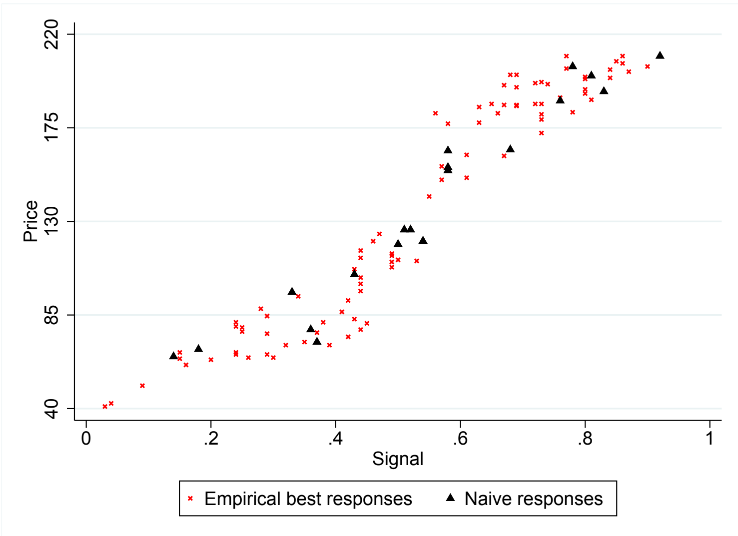

Figure 1.4: Sells and buys within the relevant area in treatment SEQ.

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

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

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

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

Hypothesis 2. Hypothesis 2 compares the buy and sell decisions of traders 1 and 2 in treatment SIM. Figure 1.5 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.[15]

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

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

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

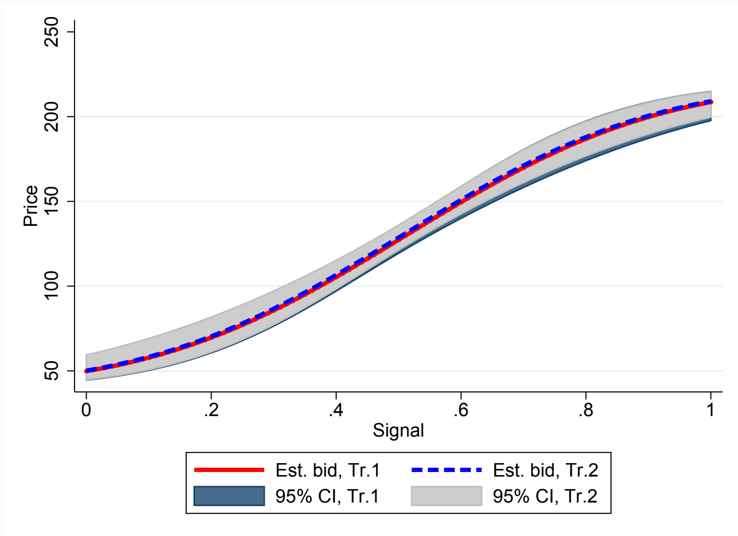

Figure 1.5: Estimated average bids of traders 1 and 2 in treatment SIM.

1.5 Possible drivers of information neglect

1.5.1 Signal strength

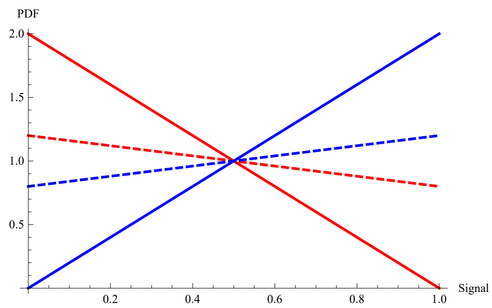

Figure 1.6: Signal distributions for trader 1 (solid) and trader 2 (dashed) in LSQ treatments.

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

We examine the hypothesis by introducing an asymmetric signal strength between trader 1 and trader 2, keeping the rest of the design unchanged. In two additional treatments with "Low Signal Quality", LSQ-SIM and LSQ-SEQ (with $N = 70$ and $N = 68$, respectively), trader 2's signal is less informative. The densities in the new treatments are depicted in Figure 1.6 and take the following form.

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

Behavior of trader 2 deviates from the naive prediction in both treatments LSQ-SIM and LSQ-SEQ. Trader 2s react to their signals more strongly than predicted by naive bidding (see Figure A1). A comparison with the bids in the main treatments SIM and SEQ thus supports the conjecture that subjects pay more attention to market information when they are less informed privately.

However, the discrepancy between the two market mechanisms increases with information asymmetry.

The share of naive decisions in treatment LSQ-SEQ (22%, black triangles in Figure 1.7b) is much smaller than in LSQ-SIM (44%, black triangles in Figure 1.7a). This significant difference ($p = 0.0003$, Wald test) corresponds to a steeper estimate of the average bidding curve in LSQ-SEQ, see Appendix Figure A1. Tables A1 to A3 in the appendix also show that differences in frequencies of buys and sells between the two mechanisms are highly significant for various signal ranges, and that they tend to be larger than in the comparison of SIM and SEQ. For example, participants in the role of trader 2 of LSQ-SEQ act very frequently against their own signal. In sum, the importance of trading mechanisms for rational decision making prevails under the new informational conditions.

[16] We thank an anonymous referee for raising this hypothesis.

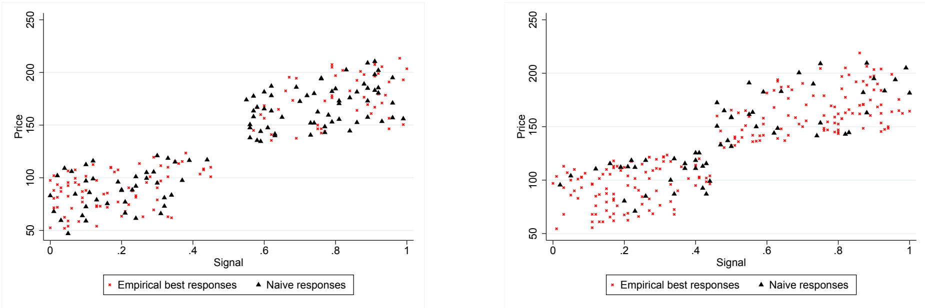

(a) Trader 1

(b) Trader 2

Figure 1.7: Buys and sells consistent with either naive bidding or EBR in treatments LSQ-SIM and LSQ-SEQ, respectively.

1.5.2 Strategic uncertainty

Strategic uncertainty adds to the complexity of the trading game. For an accurate interpretation of price, participants in the role of trader 2 need to consider the trading behavior of trader 1 and their ability to do so may vary between simultaneous and sequential mechanisms. In other words, the necessity to assess the human-driven EBR (not just the simpler BN response) may lead to less optimal behavior by trader 2 in treatment SIM relative to SEQ.

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

For better comparison with the main treatments, the instructions of the NP1 treatments retain not only much of the wording but also the chronological structure of the main treatments. Participants in NP1 treatments thus learn about the existence of p_1, which is presented to them as a random "initial value" of the asset's price, and they learn that the market maker observes an additional

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

signal that is correlated with the asset's value. Like in the main treatments, the instructions display the updating rule (1.3) and explain that it results in the price p_2 at which the participants can trade and which reflects the expectation of the asset's value, conditional on the market maker's additional signal but not conditional on the participants' own signal.

The data show no strong differences between the NP1 treatments and the main treatments. Appendix Figure A2 shows that the estimated bidding curve in NP1-SIM exhibits the same slope as the curve in SIM, with a mild downward shift, whereas behavior in NP1-SEQ is very close to that of SEQ.[18]

Most notably, the effect of simultaneous versus sequential trading 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 (see Tables A1 and A2 in Appendix A.1). Figure 8 shows the individual decisions for cases where naive and rational predictions differ, in treatments NP1-SIM and NP1-SEQ, respectively.

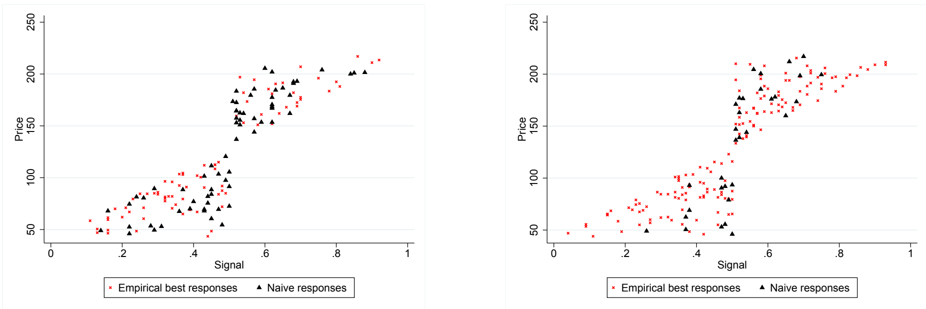

(a) NP1-SIM **(b)** NP1-SEQ

Figure 1.8: Naive v Bayesian in NP1-SIM and NP1-SEQ.

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

1.5.3 Number of decisions per treatment

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

We therefore introduce a "hypothetical" sequential treatment (Hyp-SEQ) with 62 participants, which rules out higher dimensionality of strategies as a source of difficulty. Treatment Hyp-SEQ is analogous to SEQ in that after learning trader 1's price p_1, participants in the role of trader 2 specify their buying or selling preferences for only a single price $\hat{p}_2$. However, $\hat{p}_2$ is only a candidate price as $\hat{p}_2$ is equiprobably drawn from the two price values that are possible after updating via rule (1.3). Participants decide whether they would buy or sell at $\hat{p}_2$ and the decision is implemented if and only if trader 1's demand induces the realization $p_2 = \hat{p}_2$. Otherwise, trader 2 does not trade and makes zero profit.

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

Average bidding shows no large difference between treatments SIM and Hyp-SEQ, or between traders 1 and 2 of treatment Hyp-SEQ: The estimated bid functions in Appendix Figures A3a and A3b exhibit approximately the same slope. Moreover, the Appendix A.2 also shows that naive bidding cannot be rejected for treatment Hyp-SEQ, in multiple binomial testing.

However, Figure 1.9 and Table A4 in the Appendix show that the frequency of making suboptimal decisions (η) in Hyp-SEQ lies well in between those of SEQ and SIM. The significant difference between treatments SIM and Hyp-SEQ (0.38 versus 0.28, p=0.022, one-sided t test) shows that reducing the set of hypothetical prices considerably improves decision-making. Yet, the frequency of naive

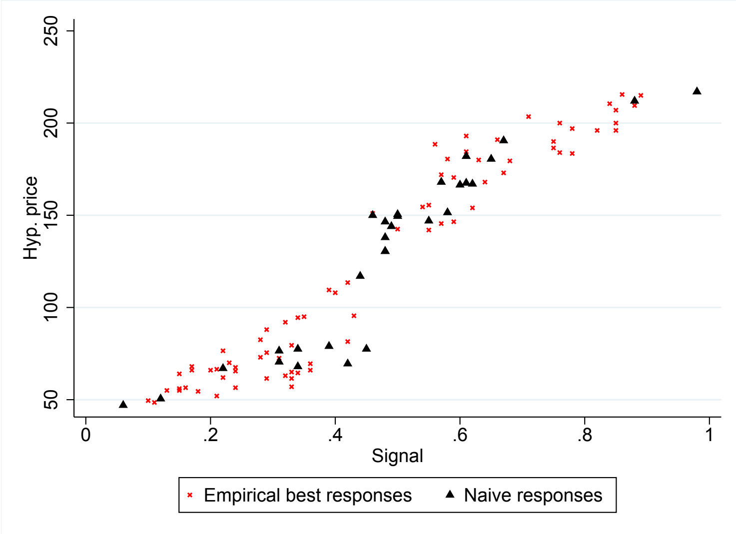

Figure 1.9: Naive v EBR in Hyp-SEQ.

decisions is still significantly higher in Hyp-SEQ than in the fully sequential treatment SEQ, (0.28 versus 0.19, p=0.081, one-sided t test).[19] Altogether, we conclude from the above tests that reducing the number of hypothetical trading decisions reduces the degree of naiveté, but does not eliminate it.[20]

1.5.4 Random utility model

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

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

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

of the available pieces of information:

$$P(X_2 = 1|u_i, s_2, p_2) = \frac{e^{\lambda(\hat{E}[\theta|p_2,s_2]-p_2+u_i)}}{1 + e^{\lambda(\hat{E}[\theta|p_2,s_2]-p_2+u_i)}} \tag{1.9}$$

with

$$\hat{E}[\theta|p_2, s_2] = 40 + 180 \cdot \hat{P}(\theta = 220|p_2, s_2) \tag{1.10}$$

$$\hat{P}(\theta = 220|p_2, s_2) = [1 + LR(s_2)^{-\beta} \cdot LR(p_2)^{-\alpha}]^{-1} \tag{1.11}$$

The choice probability (1.9) depends on subjectively expected payoff, $\hat{E}[\theta|p_2, s_2] - p_2$. The parameter λ reflects the precision of the logistic response and u_i is the random utility shifter, which we assume to be normally distributed with mean 0 and variance σ_u^2. To allow for irrational weighting of information, we introduce the subjective posterior probability of the event that $\theta = 220$, given by $\hat{P}(\theta = 220|p_2, s_2)$. Analogous to the method introduced by Grether (1992), we let the posterior probability depend on the likelihood ratios of the signal and the price, $LR(s_2) \equiv \frac{P(\theta=220|s_2)}{P(\theta=40|s_2)}$ and $LR(p_2) \equiv \frac{P(\theta=220|p_2)}{P(\theta=40|p_2)}$, respectively. The likelihood ratios are exponentiated by the potentially irrational weights β and α that the participant assigns to the signal's and the price's informational content. A participant with naive beliefs (a 'fully cursed' participant) would correctly weight the signal, $\beta = 1$, but would ignore the information in the price, $\alpha = 0$. An intermediary level of cursedness translates into α between 0 and 1. A rational trader would correctly weight the signal and the price, $\beta = \alpha = 1$. The model also allows for an over-weighting of the signal or the price, by letting β or α exceed 1.

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

The estimates are reported in Table 1.1 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 A4). Traders' 2 weighting of

Table 1.1: 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: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. 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 private signal decreases from 2.54 to 1.36 between the simultaneous and the sequential treatments. Both of these β estimates significantly differ from 1, but in the sequential treatments β lies significantly below trader 1's weighting of the private signal ($p = 0.0298$, Wald test).
In the simultaneous mechanisms, the estimated α of 0.60 lies well below the optimal value 1, albeit at a somewhat marginal statistical significance of $p = 0.09$. While this difference from 1 reflects the hypothesis that participants pay too little attention to the price's informativeness, we can also reject the extreme formulation of Hypothesis 2, stating that participants are fully naive: α differs significantly from 0.
In the treatments with sequential mechanisms, the perceived levels of informativeness of signal relative to price are reversed. These treatments induce a significant over-weighting of the price's likelihood ratio ($\alpha = 1.85$).[21] Overall, the evidence from sequential treatments shows that the prior distribution of θ is under-weighted and that, confirming Hypothesis 1, sequential markets reveal a significantly stronger inference from the price than simultaneous markets.

1.6 Discussion: Information neglect in markets

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

Pricing. A natural measure of price efficiency is the speed at which price aggregates 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

[21] 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.

the implications of naiveté on price are by Hong and Stein (1999) and Eyster et al. (2015). They both find, with different models, that the presence of naive traders creates a bias of prices leaning towards their ex-ante expectation. The reason is that naive traders are likely to engage in excessive speculation based on their own signal—they bet against the market price too often. This pushes price towards its ex-ante mean.[22]

Testing this implication requires the simulation of a specific price mechanism after trader 2 has completed her trades. For simplicity and for consistency with the rule governing p_2, we calculate the price that a market maker would 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 realized demand of traders 1 and 2, assumed to follow the Bayes-Nash prediction. In our main treatments SIM and SEQ, the price for a hypothetical trader 3 is thus a simple function of p_2 and x_2:[23]

$$p_3 = \begin{cases} \frac{-8800+310p_2}{50+p_2} & if\ x_2 = 1 \\ \\ \frac{-8800+50p_2}{310-p_2} & if\ x_2 = -1 \end{cases}$$

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 either both traders buy or both traders sell, in Bayes Nash Equilibrium. For all cases where s_1 and s_2 lie on the same side of 0.5, Figure 1.10a shows the resulting distribution of Bayes Nash price p_3 as a dotted line, with much probability mass located towards the extremes. In contrast, if trader 2 bids naively, then she will tend to sell at high prices and buy at low prices, creating excessive density of p_3 near the center of the distribution (light grey line).

Figure 1.10a 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

[22] 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. Eyster et al.'s (2015) model uses partially cursed equilibrium to show the bias in pricing, using a more standard (and more static) model of financial markets with incomplete information akin to that in Grossman (1976).

[23] 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.

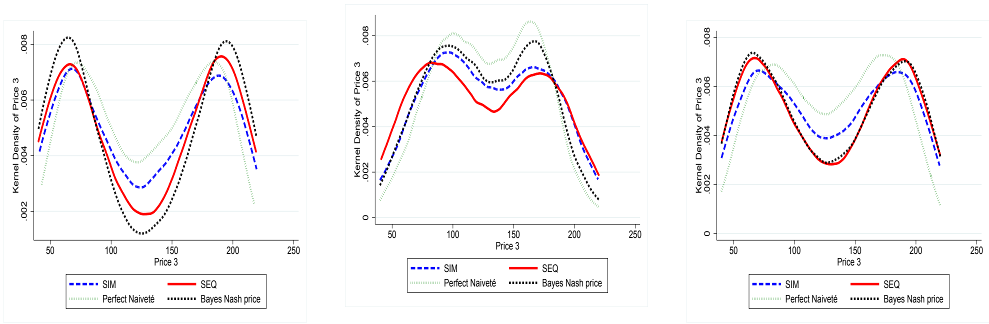

(a) s_1, s_2 in same direction

(b) s_1, s_2 in opposite direction

(c) All signals

Figure 1.10: Kernel density of efficient price 3 after naive, rational and actual demand of traders 1 and 2 in SIM and SEQ .

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 1.10b 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 1.10c depicts the densities when taking into account all observations. Overall, the price distribution in treatment SEQ has a more pronounced bi-modal shape.

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

To quantitatively assess price efficiency under the two treatments, we ask about the variance of fundamental value conditional on the price, $Var[\theta|p_3]$. It captures the error in market expectations given information contained in p_3. Conditional variance is significantly lower in treatment SEQ than in SIM, at high level of significance (p=0.00, nonparametric median test, taking each market as a unit of observation) and with a somewhat sizable difference: in treatment SIM, the price explains on average 21% of the variance in the asset value, versus 27% in treatment SEQ.[24]

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

Profits. The difference between simultaneous and sequential mechanisms also affects the distribution of profits of trader 2. A corresponding difference occurs in each of the relevant treatment comparisons, but it is economically small (our experiments were not designed to generate big payoff differences between treatments) and is statistically significant only in the comparison LSQ-SIM versus LSQ-SEQ, i.e. with asymmetry in the informativeness of signals. Less informed traders benefit from sequential information processing, where the employed updating is more rational. The results on mean and median profits of each treatment is in Table A5 in the Appendix. It is also noteworthy that the distribution of profits conditional on price p_2 in LSQ-SEQ is mirror-inverted to the one in LSQ-SIM (see Figure A5b): the majority of traders in LSQ-SIM lose significant amounts, whereas the majority of traders in LSQ-SEQ make gains. This hints at the importance of pre-trade transparency to restrain insider trading in real-world markets. Naive later traders may suffer if they are poorly informed.

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

1.7 Conclusion

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

2 Trading under ambiguity and the effect of learning

2.1 Introduction

The link between ambiguity and trading activity is not obvious. Ambiguity is a fundamental feature of financial markets, where objective probabilities of the states of nature are often unknown. The extent to which traders perceive ambiguity determines whether ambiguity begets more or less trading activity. Speculative trade increases, for instance, when ambiguity enables investors to develop subjective beliefs that are sufficiently heterogeneous. Several studies that find a positive correlation between trading volume and measures of uncertainty support this hypothesis (see Karpoff 1987 for an extensive survey of this literature).

In contrast, trading activity decreases if investors dislike events with unknown probabilities. In that case, investors who already have state-invariant positions will not re-allocate their portfolio. In incomplete markets, staying out of the market might even be the sole possibility to avoid ambiguous trades. Diminished willingness to trade can therefore be rationalized with ambiguity-averse preferences. Various models of ambiguity aversion (e.g. Choquet expected utility (CEU), maxmin expected utility (MEU), α-maxmin expected utility (α-MEU)) depart from expected utility theory by modeling decision-makers who consider different distributions for opposite actions: one for going long and one for going short. The ambiguity-averse seller short-sells at higher prices, the ambiguity-averse buyer displays a lower willingness to pay. In between, there is a range of prices at which buyer and seller do not agree on trade. In line with this rationale, Antoniou et al. (2015) find equity flows to be negatively related to ambiguity.

The different movements in trading volume can be reconciled if one acknowledges that only ambiguity perceived as such can impede trading (Dimmock et al., 2016). While ambiguity-averse traders conceive a set of probability distributions, traders with subjective beliefs do not perceive any ambiguity. The extent to which traders perceive ambiguity is, however, likely to vary with incoming information. As ambiguity often triggers enhanced information acquisition, understanding traders' reactions requires first examining how they process information under ambiguity. The way investors update ambiguous beliefs upon incoming information defines their perception of ambiguity and, hence, their willingness to trade.

This paper offers a systematic comparison of willingness to trade assets with ambiguous and unambiguous return distributions, in a stylized incomplete market with one uncertain asset and money. To examine the relevance of learning, the experiment studies investment decisions under ambiguity across two information conditions: one where investors base their decisions on given probabilities; and a second where investors receive additional information before investing. Two main updating rules serve as theoretical benchmarks: *full Bayesian updating* (henceforth FBU, Jaffray 1989; Pires 2002) and *maximum likelihood updating* (henceforth MLU, Gilboa and Schmeidler 1993). With FBU, subjects update a set of priors prior by prior and evaluate the resulting set of posteriors according to their ambiguity preferences. With MLU, on the other hand, subjects consider a subset of priors that maximizes *ex-ante* the probability of receiving the observed information. Additional information leads an agent to discard unlikely priors and to perceive substantially less ambiguity. Eventually, he will come to a single posterior belief and will not perceive any ambiguity at all. In this way, the arrival of information may generate a singleton posterior, which depends on the nature of the arrived information and, therefore, might be heterogeneous across agents.

The experiment is designed as an individual decision-making environment where subjects cannot learn from the market. In addition, any risk-sharing motive to trade is excluded because subjects start with a riskless position. A 2x2 design allows comparing decisions across two dimensions. The first dimension varies the degree of uncertainty by comparing decisions under risk versus ambiguity. The second dimension distinguishes between situations where information about return distributions is released at once and those where information is processed sequentially. The design is implemented with 2 treatments, such that the first dimension of variation is analyzed in a within-subject comparison and the second dimension between subjects. Treatment "No Learning" (**NL**) investigates the relation between ambiguity and investment decisions when belief updating is not required. Treatment "Learning" (**L**) examines the ambiguity effect when traders receive additional information after prior distributions are specified.

In the two treatments, participants submit a bid and an ask quote for an uncertain asset. In some rounds, participants learn the objective probability distribution of the asset's value and, thus, invest in a risky asset. In other rounds, they receive imprecise information about the distribution, which makes the latter ambiguous. While in treatment NL information about the distribution is revealed at once, participants in treatment L learn the distribution across two stages: They first receive information about a prior distribution, and observe then an additional signal.

One main result is that participants express a lower willingness to trade by choosing significantly wider bid-ask spreads when returns have ambiguous distributions. The decrease in market participation even persists when subjects learn distributions progressively. This result adds to the evidence of ambiguity aversion found in a multitude of Ellsberg experiments (i.a. Chow and Sarin 2002; Halevy 2007, and Camerer and Weber 1992 for a review of the literature). It shows that ambiguity aversion manifests itself in spreads when portfolio reallocation is not possible. The average ambiguity premium in long and short positions amounts to 20% and 16.4% of the expected value, respectively, and is in line with previous findings (Yates and Zukowski 1976; Bernasconi and Loomes 1992 and the references in Camerer and Weber 1992). The ambiguity premium over and above the risk premium cuts down trade by, on average, 12 percentage points, and mean profits by 30%. These findings confirm that ambiguity aversion is well suited to model freezes in trading activity.

A second main result is that learning generates more extreme quotes. Yet, there is no evidence of subjects being predominantly MLU agents. MLU predicts small to zero spreads, but subjects choose the same average spread when the same ambiguous distribution is learned progressively. The evidence in favor of FBU, too, is limited: Bids and asks are significantly lower (higher) after the arrival of a low (high) signal. The chosen quotes are rather consistent with updating second-order beliefs on ambiguous probabilities. A bulk of 36.87% decisions for ambiguous prospects is centered around Bayesian updates of the mid-prior. The remainder of quotes discloses heterogeneity in the way of updating ambiguous beliefs. One noticeable group is insensitive to additional information and refrains from trading. Another major group consists of extreme updaters who choose to trade at all prices.

In sum, the results identify a negative relation between ambiguity and willingness to trade that is robust to the information condition. The relation between ambiguity and trading volume, though, is not conclusive since a lower willingness to trade does not directly translate into lower trading volume. Despite ambiguity-averse trading preferences, differential information combined with Bayesian updating of recursive preferences may generate updated beliefs that are divergent enough to spawn speculative trading. This effect of learning matters, in particular, because the link between ambiguity and liquidity is characterized by feedback effects. Traders who do not perceive ambiguity nurture liquidity, which, in turn, encourages price discovery. In contrast, whenever ambiguity engenders a drop in liquidity, prices fail to aggregate information and are more prone to excess volatility (Dow and Werlang, 1992b; Guidolin and Rinaldi, 2010). In that case, ambiguity produces market frictions that may persist over longer periods. The experimental data suggest that gradual information processing can mitigate ambiguity effects if quotes become sufficiently heterogeneous.

Moreover, subjects' more extreme reactions with gradual information release have direct implications for discretionary disclosure policy. Miller (2002) and Kothari et al. (2009) find evidence for an asymmetric disclosure of bad and good news: While managers disclose good news immediately, they accumulate bad news before releasing them. The experimental findings indicate that the asymmetric disclosure has effects beyond the one of supporting managers' careers: It possibly dampens negative, but fosters positive stock price reactions.

This paper relates two strands of research. One strand examines the effect of ambiguity on market parameters. In theoretical models of market microstructure, for instance in Cao et al. (2005); Ui (2011), and Easley and Hara (2010), ambiguity aversion is used to model limited market participation. In line with this theoretical literature, this paper uses a stylized decision experiment to test the hypothesis that market participation decreases with ambiguity. Two other experimental studies analyze the effects of ambiguity on financial decisions. Ahn et al. (2014) individual-decision experiment confirms the heterogeneity in ambiguity attitudes, providing evidence for subjective expected utility (SEU), ambiguity aversion as well as for pessimism. Bossaerts et al. (2010) show in their market experiment that heterogeneity in ambiguity attitudes does not only affect portfolio choices, but also asset prices. Standard price predictions do not hold and prices do not aggregate beliefs if only the least ambiguity-averse traders provide the total supply of ambiguous assets. In contrast, the design in the present experiment identifies ambiguity aversion not through portfolio allocation but with chosen spreads. It focuses on individual willingness to trade and, thus, extends the study of ambiguity aversion to markets that do not provide the opportunity to fully insure against ambiguous states. A related study is Sarin and Weber (1993). They find bids and the resulting market prices for ambiguous assets to be consistently lower in sealed-bid and oral double auctions, although ambiguous and unambiguous assets had identical expected payoffs.[1] As they conclude, subjects are less willing to pay for ambiguous assets that they apparently consider as more risky. Another related work is the experimental study of Eisenberger and Weber (1995). They find no interaction between ambiguity and the buying/selling price ratio. As their focus lies on the buying/selling price ratio, willingness to pay and willingness to accept are elicited from different default positions. This study, in contrast, focuses on the individual willingness to trade by keeping the starting position constant and state-invariant. This allows for testing the prediction made in Dow and Werlang (1992a) under varying conditions.

Another strand of the literature analyzes belief updating under ambiguity. The current work contrasts from Epstein and Schneider (2008), which models updating of ambiguous information. Instead, this research evaluates belief updating of ambiguous priors when information is precise. Cohen et al. (2000) use in this context a dynamic extension of the Ellsberg experiment to differentiate between

[1] Note, in their oral double auctions subjects are endowed with assets. In that case, ambiguity-averse traders want to get rid of their uncertain endowment and drive down the offer price.

FBU and MLU behavior. They, too, find heterogeneity in updating behavior. The behavior of a non-negligible amount of subjects is consistent with MLU, but FBU seems to be the more predominant updating rule in their implementation of the Ellsberg-experiment. The current paper emphasizes the importance of these two updating rules for trading activity and provides another framework to distinguish between them, and possibly more updating rules. The relevance of alternative or modified updating rules is shown in De Filippis et al.'s (2016) experiment with both social learning and private signals, where subjects' updated beliefs are more consistent with *likelihood ratio test updating*, a generalization of MLU. Another related experiment also studies learning in ambiguous asset markets: Baillon et al. (2013) investigate learning with a natural source of uncertainty. In their individual decision-making design, subjects submitted ask prices for options on initial public offerings (IPOs). Using the neo-additive model (Chateauneuf et al., 2007), they find no evidence for pessimism (ambiguity aversion). Furthermore, whereas pessimism is not affected by the arrival of new information, sufficient information reduces likelihood insensitivity. The following experiment adds to this literature and contrasts markets with ambiguity shocks and ambiguous markets with gradual information release. Moreover, it compares learning in ambiguous markets to learning in risky markets to identify learning effects that are specific to ambiguity.

The paper is organized as follows. Section 2.2 presents the stylized decision model and the theoretical predictions. Section 2.3 describes the implementation in the experiment. The results are presented in Section 2.4. Section 2.5 discusses their implications and concludes.

2.2 The theoretical framework

2.2.1 Investing in ambiguous versus risky prospects

A stylized decision problem

Consider a simple investment opportunity in a market with two states and one risky asset. The investor may invest in *one unit* of the risky asset with value $V \in \{V_L, V_H\}$. The probability for the high-value state corresponds to

$Pr(V = V_H) =: \pi$.

The investor is endowed with cash W_0 and tenders both a bid quote, b, and an ask quote, a, before knowing the transaction price, p. The price p is exogenous and is drawn from a uniform distribution, i.e. $p \sim U[V_L, V_H]$. The agent's demand corresponds to:

$$X = \begin{cases} +1 & if\ p \leq b \\ -1 & if\ p \geq a \\ 0 & otherwise \end{cases}$$

The agent is a price taker: At the end, he will pay a price p that he cannot influence and that will possibly differ from his quotes b and a. The quotes b and a merely determine the probability that a buy or a short-sale (henceforth sell) occurs. A higher bid b, for instance, increases the probability to buy, as the random price p is more likely to fall below it. Notice that the agent will always trade whenever the bid equals the ask. The investor's wealth at the end of the period is $W_1 = W_0 + (V - p) \cdot X$.

Predictions

Denote Π^* as the agent's subjective set of beliefs about π, the probability for the high-value state.

For the benchmark analysis of expected utility, assume that the agent holds a single probability belief π, i.e. Π^* is a singleton. Under risk-neutrality, he buys at prices below his expected valuation, sells at prices above it, and therefore sets $a^* = b^* = E[V]$. A risk-averse agent, on the other hand, chooses a strictly positive spread between bid and ask, with $b^* < E[V]$ and $a^* > E[V]$ (the simple proof is in the Appendix B.2).

Optimal values of bid and ask may change when the agent perceives ambiguity about π. That is, if he contemplates an interval of probabilities $\Pi^* = [\pi_l, \pi_h]$, bid and ask quotes adjust to his ambiguity preferences. Different models of ambiguity aversion will then predict different trading quotes. The present argumentation follows Dow and Werlang (1992a), but uses the intuitive model

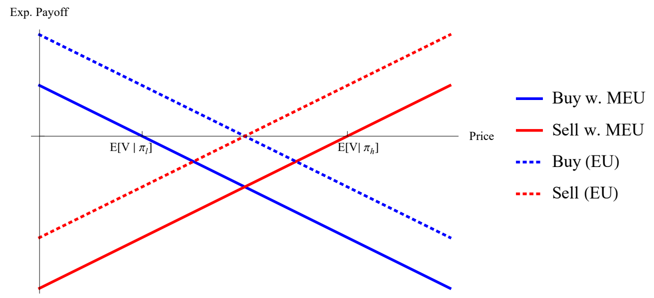

Figure 2.1: Expected payoff of a buy and a sell as a function of the price for risk-neutral EU (dashed lines) and MEU (solid lines) agents.

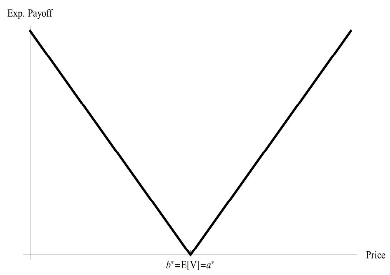

(a) EU

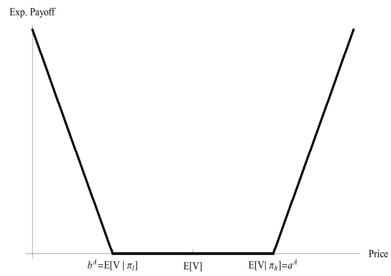

(b) MEU

Figure 2.2: Expected payoff with optimal strategy of risk-neutral (a) EU and (b) MEU agent.

of maxmin expected utility (MEU - Gilboa and Schmeidler 1989) instead of Choquet expected utility.
An MEU agent evaluates different actions with different probability distributions. He considers the worst possible expected outcome, which differs between the cases where he buys and sells. A risk-neutral MEU agent buys if

$$p \leq \min_{\forall \pi \in [\pi_l, \pi_h]} E[V|\pi].$$

He sells if

$$p \geq \max_{\forall \pi \in [\pi_l, \pi_h]} E[V|\pi].$$

The expected payoff functions of ambiguity-averse buying and selling strategies are shifted downwards, relative to the case of expected utility (see Figure 2.1). Due to the fact that willingness to buy and willingness to sell do not intersect at a single strictly positive price, there is a region of prices at which zero holding of the asset is optimal (cf. Dow and Werlang 1992a, see Figure 2.2).[2] Given risk-neutrality, models of ambiguity aversion with kinked preferences predict wider spreads for ambiguous than for unambiguous prospects. Predictions under risk aversion, however, depend on the preferences and can vary widely. Importantly, for smooth ambiguity preferences like those studied in Klibanoff et al. (2005) (henceforth KMM), the spread converges to the spread of an expected utility maximizer when ambiguity aversion converges to neutrality. The validity of smooth preferences is discussed in Section 2.4.3. The main objective of the experiment is not to identify kinked from smooth preferences, but to generally compare spreads for ambiguous and unambiguous assets. Differences in spreads are used to test whether ambiguity leads to a premium that is, on average, larger than the risk premium.

2.2.2 Introducing information

Consider now an environment where the agent receives an informative signal prior to investing. The signal $s \in \{\vartheta_L, \vartheta_H\}$ is binary, symmetric and correct with probability $q = P(s = \vartheta_L | V = V_L) = P(s = \vartheta_H | V = V_H)$. Henceforth, the prior and posterior beliefs are denoted with $Pr(V = V_H) =: \mu$ and $Pr(V = V_H | s, \mu) =: \rho$, respectively.

For exposition, predictions are presented for risk-neutral EU and MEU agents. The difference in predictions holds under risk aversion as well.

Bayesian updating

A rational agent who has a single prior belief μ applies Bayes rule, then quotes a bid and an ask $b = a = E[V|s]$. That is, the risk-neutral EU agent adjusts the quotes to information, but holds a zero spread before and after information. A

[2]When the starting position is risky instead of riskless, the general result holds as long as the returns of risky and ambiguous assets are negatively correlated. The possibility to hedge the ambiguous asset with the risky one decreases the range of non-participation, but does not fully eliminate it (Epstein and Schneider, 2010).

risk-averse EU agent holds the same non-zero spread for the same belief value, regardless of the belief being a prior or a posterior belief.

In contrast, if the prior is ambiguous, optimal quotes depend on the way the agent updates ambiguous beliefs. It is still an open question how agents update ambiguous beliefs. The literature has proposed various updating rules (see Epstein and Schneider 2007, Gilboa and Schmeidler 1989; Hanany and Klibanoff 2007; Jaffray 1989; Klibanoff et al. 2009). Here we focus on two main concepts that do not require any specific preference model. Moreover, the two paradigms make maximum opposite predictions with respect to the spread.

Full Bayesian updating

Agents with multiple priors apply FBU when they update prior by prior to end up with a set of posteriors. In the case where an agent considers solely the support of prior probabilities (without having second-order beliefs over priors), he will update the two extreme priors to two extreme posteriors. Therefore, unless $q = 1$, FBU does not fully eliminate ambiguity. The choice of the relevant posterior and hence the evaluation of an action depend then on ambiguity preferences. For instance, an MEU agent with a high signal ($s = \vartheta_H$) buys an asset if

$$p \leq \min_{\mu \in [\mu_l, \mu_h]} E[V|s = \vartheta_H, \mu].$$

He therefore bids $b = E[V|s = \vartheta_H, \mu_l]$. Analogously, his ask corresponds to $a = E[V|s = \vartheta_H, \mu_h]$, with $b < E[V|s = \vartheta_H, \frac{\mu_l + \mu_h}{2}] < a$.
Hence, the ambiguity-averse trader chooses a non-zero spread both before and after the updating. Its value depends on Π^*, the set of probabilities that the trader considers as possible.

Maximum likelihood updating

With MLU, the information received pins down the set of priors that will be updated. The prior that has *ex-ante* the highest probability to generate the informational event is given *ex-post* the highest likelihood. In our specific setting, an agent observing a high signal ($s = \vartheta_H$) assigns the highest likelihood to the highest prior μ_h. The agent therefore postulates a single posterior whenever a single prior maximizes the likelihood of having generated the informative event. If so, the signal completely eliminates the perception of ambiguity. The agent adjusts his belief to one of the two extremes, depending on the signal being high or low.

The optimal bid satisfies then:

$$p \leq E[V|\mu^*, s] \quad \text{with } \mu^* = \arg \max_{\mu \in [\mu_l, \mu_h]} \ell(\mu|s),$$

where $\ell(\mu)$ represents the likelihood of a prior. The same prior μ^* satisfies the likelihood in the condition for the optimal ask:

$$p \geq E[V|\mu^*, s] \quad \text{with } \mu^* = \arg \max_{\mu \in [\mu_l, \mu_h]} \ell(\mu|s),$$

Hence, a risk-neutral MLU trader with $(s = \vartheta_H)$ and a unique posterior belief $\rho(\mu^*, s = \vartheta_h)$ chooses equal bid and ask $b = a = E[V|\mu^*, s = \vartheta_H]$.

A fundamental difference between FBU and MLU in this setting is, therefore, that the ranking of states is determined by different factors. When an agent applies FBU, the ranking of states depends on his ambiguity preferences and is determined by the long or short position (Mukerji and Tallon, 2001). An agent using MLU ranks the states according to his information.

2.2.3 Hypothesis and treatment effect

As shown in Section 2.2.1, under the assumption of risk-neutrality, ambiguity aversion introduces a bid-ask spread. In the case of risk-averse preferences, ambiguity aversion leads to wider spreads than the spread chosen at the mid-probability. Furthermore, the analysis of ambiguity aversion goes beyond any spread increase that can be explained with subjective expected utility. Consider, for instance, an ambiguous set of probabilities $[\pi_l, \pi_h]$ that encompasses the probability $\pi = .50$, at which theory predicts a maximum spread with risk-averse utility functions. If the mid-probability of the set differs from 50% (i.e. $(\frac{\pi_l+\pi_h}{2}) \neq .50$), a subjective belief of $\Pi^* = .50$ can rationalize a wider spread than the spread chosen at the mid-probability. In contrast, subjective beliefs fail to rationalize spreads that are wider than any chosen spread at every unambiguous probability $\pi \in [\pi_l, \pi_h]$. In this context, the experiment targets evidence in favor of ambiguity aversion that cannot be simultaneously rationalized by subjective expected utility.

Hypothesis 3 *Ambiguous probabilities induce wider bid-ask spreads than unambiguous probabilities:*

$$E[a - b|\pi \in [\pi_l, \pi_h]] > E[a - b|\pi], \quad \forall\, \pi \in [\pi_l, \pi_h]. \tag{2.1}$$

Therefore, bid-ask pairs for an ambiguous set $[\pi_l, \pi_h]$ that are more divergent than bid-ask pairs chosen at any $\pi \in [\pi_l, \pi_h]$, i.e. at all unambiguous probability values in the same set, are interpreted as evidence in favor of ambiguity aversion.

If subjects are ambiguity-averse, changes in their perception of ambiguity can translate in variation of the spread. In a second step, differences in quotes are used to assess how gradual information processing affects the perception of ambiguity.

The experiment is designed such that full Bayesian updaters would quote the same bid-ask pairs for ambiguous prospects in the two treatments NL and L. In contrast, maximum likelihood updaters would perceive substantially less ambiguity and choose smaller spreads in treatment L. To this effect, the comparison across treatments focuses on rounds with identical sets of marginal and FBU probabilities. Identical spreads in the two treatments indicate that subjects perceive the same support of probabilities, which would provide evidence in favor of FBU:

$$\texttt{Under FBU:}\ E[a - b|\rho \in [\rho_l^{FBU}, \rho_h^{FBU}]] = E[a - b|\pi \in [\pi_l, \pi_h]]$$
$$\texttt{with}\ [\rho_l^{FBU}, \rho_h^{FBU}] = [\pi_l, \pi_h].$$

However, smaller observed spreads in treatment L are more consistent with beliefs resulting from MLU than from FBU, suggesting that subjects react more strongly to information and perceive less ambiguity than an FBU agent would:

$$\texttt{Under MLU:}\ E[a-b|\rho \in [\rho_l^{FBU}, \rho_h^{FBU}]] < E[a-b|\pi \in [\pi_l, \pi_h]]$$
$$\texttt{with}\ [\rho_l^{FBU}, \rho_h^{FBU}] = [\pi_l, \pi_h].$$

Thus, comparing the average spread between treatments NL and L for the same support of marginal and FBU probabilities allows for differentiating between FBU or MLU.

2.3 Experimental design

2.3.1 Treatment No Learning (NL)

Treatment **NL** consists of 20 rounds. In each round, subjects start with an endowment of cash W_0 and tender both a bid and an ask ($b, a \in [V_L, V_H]$, $b < a$).[3] At the beginning of each round, subjects receive information about the uncertainty of the investment. At that stage, they learn whether π is ambiguous or not. The uncertainty in the asset's value is visualized by displaying "urn A" that contains 100 balls in a mixture of red and blue balls. To determine the asset value, the computer draws a ball (henceforth "value ball") from urn A: If a red ball is drawn, the asset takes the value V_L. The asset takes the value V_H, if the value ball is blue.

The proportion of red and blue balls in urn A varies across rounds (see Table 2.1 for the chosen parameters) and is shown to the subjects. That is, subjects learn π for risky prospects by observing the exact number of red and blue balls in urn A. When the distribution is ambiguous, the exact proportion of red and blue balls is not disclosed: Instead, subjects observe a minimum number of red and a minimum number of blue balls. The remaining balls in urn A are depicted as grey. Thus, subjects learn an interval range for π (e.g. $\pi \in [.15, .85]$), but they do not know its exact value (see Figures B1 in Appendix B.1 for examples of urn A with unambiguous and ambiguous distributions).

To implement payoffs in ambiguous rounds, the computer chooses with equal probability a value in $[\pi_l, \pi_h]$. Subjects, however, did not receive any information

[3] The submission of two separate quotes allows subjects to reflect on a buy and a sell separately, as presumed in models with kinked preferences.

about how the true composition of urn A is determined when π is ambiguous. Subjects quote then bid and ask on a second, separate screen.

2.3.2 Treatment Learning (L)

Treatment **L** is almost identical to treatment NL, except that it contains an interim, second stage in which subjects are given an additional signal about the asset value.

In the first stage, subjects receive information about the prior μ. Like the subjects in treatment NL, they observe the composition of urn A, which is ambiguous or unambiguous, depending on the round of the experiment.

In a second stage, they receive an additional signal. They observe the color of another ball (henceforth “signal ball") that is drawn from a second urn. The choice of the second urn sets the correlation between the signal and the asset value: If the value ball is red, i.e. the asset has value V_L, the signal ball is drawn from “urn L” that consists of 75 pink and 25 green balls. If the value ball is blue, the signal is drawn from “urn H” that consists, in turn, of 75 green and 25 pink balls. Hence, the signal is correct, i.e. a pink (green) ball is drawn when the value ball is red (blue), with a probability of 75%.

Subjects observe the color of the signal ball (pink or green), but they do not know whether the signal ball is drawn from urn L or urn H (in other words, they do not know whether the asset has value V_L or V_H). Figure B2 in Appendix B.1 depicts an example of the screen at the second stage.

2.3.3 Experimental procedures

The computerized experiment was run in the laboratory of Technical University Berlin.[4] In total, 67 and 66 students participated in treatments NL and L, respectively. Each treatment was run with 3 sessions of ca. 22 subjects.

The trading game started once all participants read the instructions and answered an understanding test correctly. After all subjects completed the trading game, control measures of general attitudes towards risk, uncertainty and ambiguity were elicited.

[4]The experimental interface was programmed with the software z-tree (Fischbacher, 2007). Participants were recruited with the ORSEE database (Greiner, 2004).

The asset can take either the value $V_L = 0$ or $V_H = 100$. Subjects start each round with a cash endowment $W_0 = 100$.

Table 2.1: CHOSEN VALUES FOR THE PROBABILITY π AND THE PRIOR μ WITH CORRESPONDING BAYESIAN POSTERIOR ρ

	No Learning	Learning		
			$\rho(s = \vartheta_L)$	$\rho(s = \vartheta_H)$
Risk	$\pi = .05$	$\mu = .05$	$\rho = .02$	$\rho = .14$
	$\pi = .15$	$\mu = .15$	$\rho = .05$	$\rho = .35$
	$\pi = .35$	$\mu = .35$	$\rho = .15$	$\rho = .62$
	$\pi = .50$	$\mu = .50$	$\rho = .25$	$\rho = .75$
	$\pi = .65$	$\mu = .65$	$\rho = .38$	$\rho = .85$
	$\pi = .85$	$\mu = .85$	$\rho = .65$	$\rho = .95$
	$\pi = .95$	$\mu = .95$	$\rho = .86$	$\rho = .98$
	$T_R = 7 \times 2 = 14$	$T_{RI} = 7 \times 2 = 14$		
	Prior	Prior	Posterior (with FBU)	
Ambiguity	$\pi \in [.05; .65]$		$\rho(s = \vartheta_L) \in [.05; .65]$	
	$\pi \in [.15; .85]$	$\mu \in [.15; .85]$		
	$\pi \in [.35; .95]$		$\rho(s = \vartheta_H) \in [.35; .95]$	
	$T_A = 3 \times 2 = 6$	$T_{AI} = 1 \times 6 = 6$		
Total	$T_{NL} = 20$	$T_L = 20$		

Note: Subjects in treatment L are informed about the prior μ and the signal, but not about the Bayesian posterior ρ. Posterior probabilities are rounded to two decimal places. The parameter T denotes the number of rounds. Each parameter value occurs in two rounds, except for the ambiguous prior in L: The 6 ambiguous rounds start with the same set [.15, .85].

The set of possible probability values is chosen to be parsimonious to have enough observations for the comparison between treatments. Each treatment consists of 14 rounds with unambiguous as well as 6 rounds with ambiguous probabilities, amounting to 20 rounds in total. The variation in the unambiguous probabilities π and μ is identical in both treatments NL and L. The ambiguous rounds, on the other hand, differ between the two treatments: In L, the set of priors is fixed to $[.15; .85]$ (see Table 2.1). There, the variation in beliefs comes from the signal's value that implies either a low range for the set of FBU posteriors($\rho(s = \vartheta_l) \in [.05; .65]$) or a high range $\rho(s = \vartheta_h) \in [.35; .95]$). As described in Subsection 2.2.3, the two set of probabilities $[.05; .65]$ and $[.35; .95]$ in NL were chosen to equal the set of posterior beliefs under FBU in L. This enables to compare bids and asks for the same dispersion in probabilities, when information on the distribution is provided immediately or sequentially.

Within each treatment, participants made their decisions in alternating blocks of 7 consecutive risky and 3 consecutive ambiguous rounds. Within each block, probabilities were ordered in increasing or decreasing order for less confusion (Vieider et al., 2015). In one out of the three sessions (per treatment), the ordering of blocks were reversed. In addition, subjects played 4 trial rounds with different parameter values. Two of the trial rounds had ambiguous probabilities.

Decisions were incentivized with a random incentive system. To encourage subjects to consider each decision problem in isolation, the payoff-relevant round was chosen *at the beginning* of the trading game (Baillon et al., 2015). For this purpose, subjects threw a twenty-sided dice after the trial rounds, but before playing the 20 rounds. That is, they were aware that the payoff-relevant round was fixed during the experiment, but learned which round was chosen only at the end of the trading game.[5]

Earnings consist of a show up fee (5 EUR), plus two-third of the randomly drawn round in the trading game plus one-third of a randomly chosen task for the elicitation of preferences. The exchange rate was 0.13 EUR per experimental currency units (ECU). Minimum and maximum earnings were 5 EUR and 28.84 EUR, respectively. Subjects earned, on average, 19.50 EUR for approximately 100 minutes.

2.4 Results

2.4.1 Treatment NL

Decisions for risky prospects. Subjects make mostly risk-averse choices: A majority of bid-ask pairs have a non-zero spread. Since the distribution of spreads

[5] The instructions as well as the computer screen emphasized accordingly that hedging across rounds makes no sense once the payoff-relevant round is determined.

is highly right-skewed, analyses focus mainly on quantiles.[6] The median spread matches the risk of investing: It is hump-shaped in the probability, with a maximum at a probability of 50% (see Figure 2.3a). Furthermore, the spread is asymmetric around the probability, reflecting that increasing the bid (the ask) becomes more (less) risky with an increasing probability (see Figure 2.3b). Buying and selling are not equally risky as long as the low-value and high-value states are not equally probable. When the expected value is high, bidding is more risky than asking the expected value: A high bid entails the risk to pay a high price for a low-value asset, whereas a high ask price limits the risk of selling a high-value asset. The reverse holds when the expected value is low.[7] Overall, subjects choose a median spread of 5 ECU.

Table 2.2: Median and mean spread for various ranges of ambiguous and unambiguous probabilities.

π	$[5\% - 65\%]$	$[15\% - 85\%]$	$[35\% - 95\%]$	Total obs.	
		Median		Mean	Median
Risk	9	10	10	18.50(.825)	5
Amb.	20	28	20	29.23(1.464)	20
Diff.	-11***	-18***	-10**	-10.73***	-15***
N	804	804	804	1340	

Note: Median test (and two-sample test in means): *: p-value<.1,**: p-value<.05, ***: p-value<.01. Robust standard errors clustered at subject level (CRSE) in parentheses. The variable Amb. represents the indicator variable for rounds with an ambiguous probability.

Decisions for ambiguous prospects. Ambiguity about the probability reduces significantly subjects' willingness to trade. The median bid is shifted downwards, the median ask increases, leading to significantly wider spreads for ambiguous prospects (see Table B3 in Appendix B.3.1). Median spreads for prospects with

[6] Most analyses yield even more significant results for mean values.

[7] Subjects are more risk-averse in buying than in selling. I thank Marina Agranov for pointing to me that this finding is consistent with recent evidence showing that the willingness to sell is better at reflecting market beliefs, whereas willingness to buy reflects more personal preferences.

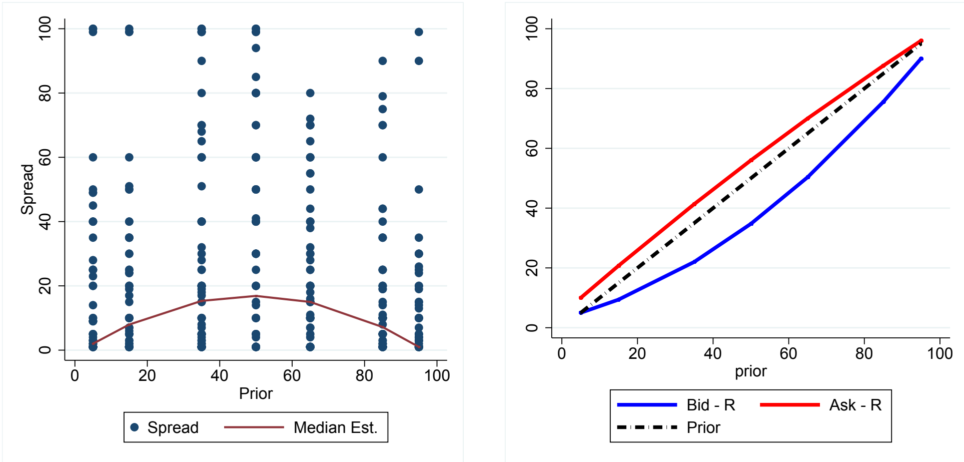

(a) Spreads and median regression **(b)** Median bid and ask.

Figure 2.3: Median spreads and quotes as a function of unambiguous priors.

ambiguous probabilities are four times as high as for unambiguous probabilities (see Table 2.2). Despite the asymmetry in risk premia for long and short positions, the ambiguity premium is almost symmetric. Subjects exhibit a median risk premium of 20% and 6.7% of the expected value in the bid and the ask, respectively. Ambiguity adds a premium of 20 and 16.4 percentage points in the bid and the ask (see Table B2 in Appendix B.3.1). In sum, Hypothesis 1 is confirmed.

Result 1 *Ambiguity in probabilities engenders wider spreads.*

As a direct consequence of the design, subjects trade and earn less when the return distribution is ambiguous. Subjects trade risky prospects in 82% of all rounds. Trades fall by 14.8% (12 percentage points) when probabilities are ambiguous. The greatest reduction of 19.3% occurs when the probability is between 15% and 85% (see Table 2.3).

The reduction in trading activity translates into significantly smaller profits. Subjects earn, on average, 41.98% (p=.0015, two-sample t-test) more in risky rounds than in ambiguous rounds (See Table B1 in Appendix B.3.1).

Table 2.3: Percentage of trades across different ranges of probabilities

π	[5% − 65%]	[15% − 85%]	[35% − 95%]	Total obs.
Risk	80.44	79.55	79.55	81.77
	(1.5)	(1.6)	(1.6)	(1.3)
Amb.	71.89	64.17	73.88	69.65
	(3.9)	(4.2)	(3.8)	(2.3)
Diff.	9.55**	15.37***	5.67	12.12***
	(4.2)	(4.4)	(4.1)	(2.6)
N	804	804	804	1340

Note: P-values of binomial test with CRSE: *: p-value<.1,**: p-value<.05, ***: p-value<.01.

2.4.2 Treatment L

This section describes first how information processing affects investment decisions. It then compares learning with ambiguous and unambiguous priors by fitting decision weights functions.

Ambiguity effects with gradual information processing. The general effects of ambiguity on the spread are robust to incoming information. In the aggregate, choices in treatment L are ambiguity-averse. Subjects choose wider spreads for ambiguous than for risky asset distributions, with increasing difference in the mean in the last 10 periods (see Table 2.4).

Yet, under ambiguity, subjects are not insensitive to information. Starting with a set of priors $\mu \in [.15, .85]$, full Bayesian inference reduces the interval of probabilities by 10 percentage points ($\Pi^*(s = \vartheta_l) = [.05, .65]$ or $\Pi^*(s = \vartheta_h) = [.35, .95]$), while MLU even eliminates ambiguity. The diminished ambiguity is expressed in subjects' quotes. The ambiguous rounds in treatment L show more trading activity than the rounds with the same set of marginal probabilities $\pi \in [.15, .85]$ in NL: The average spread for ambiguous prospects is smaller by

Table 2.4: Median and mean spread with ambiguous and unambiguous priors in Treatment L.

Rounds	1-10		11-20		1-20	.
	Med.	Mean	Med.	Mean	Med.	Mean
Risk	8.5	20.29(1.24)	10	18.79(1.13)	10	19.54(.84)
Amb.	19	24.38(1.85)	20	28.29(2.10)	20	26.33(1.40)
Diff.	-10.5***	-4.09**	-10**	-9.5***	-10***	-6.79***

Note: One-sided median test and two-sample test in means: *: p-value<.1,**: p-value<.05, ***: p-value<.01. Standard errors in parentheses.

29% (median (mean) spread of 28 (35.10) in NL vs. 20 (26.33) in L, p=.01, median test). Trading activity is higher by 22% (64% in NL vs. 79% in L, p=0.011 binomial test with CRSE). Mean profits are 34% higher (6.29 ECU more on average, p=0.088, two-sample test).

However, controlling for the range of marginal and FBU probabilities, no difference in the aggregate distribution of spreads is observable (p-value=.92 in Kolmogorov-Smirnov test, see Figures B5a and B5b in the Appendix B.3.2). Comparing rounds where marginal probabilities (π) and FBU posteriors (ρ) lie in the same interval $[.05; .65]$ discloses a small difference in the spread: Participants in NL choose a median spread of 20, whereas the median spread in L equals 15. This non-significant difference carries even less weight in the aggregate since the two treatment groups choose identical median spreads of 20 when both π and $\rho(s = \vartheta_h) \in [.35; .95]$. Apparently, subjects do not perceive substantially less ambiguity when the same information is released gradually.

Result 2 *Given the same range of marginal and FBU posterior probabilities, the aggregate distribution of spreads with ambiguous posterior beliefs does not differ from the one with ambiguous marginal beliefs.*

Therefore, data do not lend support to MLU theory. Yet, data are not completely consistent with FBU theory either: Subjects react differently to ambiguity in final probabilities than to ambiguity in posteriors. Although spreads are, in the aggregate, constant, chosen bids and asks are more extreme after information. Participants in treatment NL choose a median bid and ask of 17.5 and 50 when $\pi \in [.05; .65]$. Participants in treatment L, however, choose a median bid and ask of 10 and 40 for an FBU posterior $\rho \in [.05; .65]$ (significant differences at 5% level each). Analogously, the median bid and ask is 40 and 70.5 in the rounds where $\pi \in [.35; .95]$, but 50 and 81 in the rounds with a set of FBU posteriors $\rho \in [.35; .95]$ (significant differences at 1% level each). To examine the extent to which these quotes are compatible with either updating rule, we next consider subjects' probabilistic sophistication.

Analysis of quotes

Updating unambiguous priors. The probabilistic sophistication is analyzed with the decisions for risky prospects. First, the risky rounds in NL are used to establish a pattern between decisions and objective probabilities. Subjects should react in the same way to probabilities, regardless of probabilities being given or updated. Second, assuming that this pattern is stable - even if information is released gradually -, this pattern serves as benchmark to discuss the validity of Bayesian posterior probabilities.

The underlying regression model assesses the extent to which the bid and the ask follow the asset's expected value. Beliefs are estimated with nonlinear least squares in a seemingly unrelated regression with robust standard errors (NNLS-SUR):

$$\begin{cases} b_i = (1 - RP_b) \cdot E[V|\tilde{\tau}] + \epsilon_{i,b} \\ a_i = (1 + RP_s) \cdot E[V|\tilde{\tau}] + \epsilon_{i,s} \end{cases} \tag{2.2}$$

where $E[V|\tilde{\tau}] = V_H \cdot \tilde{\tau}$.

It is therefore assumed that bids and asks both follow the subject's expectation about the fundamental value, but potentially in a distorted way. Because subjects in treatment NL are more risk-averse in buying than in selling, the risk premium in selling RP_s is allowed to differ from the risk premium in buying RP_b. The subject's expectation is a function of his belief $\tilde{\tau}$, which does not necessarily equal the objective probability. The mapping between objective probabilities and beliefs is represented with a weighted probability function proposed by Prelec (1998):

$$\tilde{\tau}_i = e^{(-\beta(-\ln \tau)^{\alpha})}$$

The subject's belief $\tilde{\tau}$ is a weighted function of the objective probability τ. In treatment NL, $\tau = \pi$, whereas in treatment L, the objective probability is assumed to be the Bayesian posterior $\tau = \rho$.[8] The coefficient α regulates the curvature of the function. The parameter β determines the inflection point of the curve.

Table 2.5: Coefficient estimates for probability weighting function and risk premia

	NL		L	
β	0.7971	(.0576)	0.7940	(.0424)
α	0.6861	(.0612)	0.7411	(.0722)
RP_s	0.0110	(.0316)	0.0272	(.0326)
RP_b	0.2583	(.0366)	.2420	(.0280)

Note: Nonlinear least squares estimation with CRSE. Estimates are not significantly different.

The probability weighting function is in general inverse s-shaped, reflecting a general over-weighting of small and under-weighting of high probabilities. The functions do not differ between the two treatments. That is, subjects react to unambiguous, given probabilities in the same way as to unambiguous Bayesian posteriors. Assuming a stable relation between decisions and probabilities, Bayesian inference cannot be rejected.

Updating ambiguous priors. Analogous to the analysis of risky decisions, I use the data in treatment NL to establish a pattern between decisions and ambigu-

[8] An alternative definition of Bayesian inference is that subjects apply Bayes' rule to the weighted priors. As I compare subjects' reaction to objective probabilities, I use the definition of Bayesian updating that is closest to the objective probabilities.

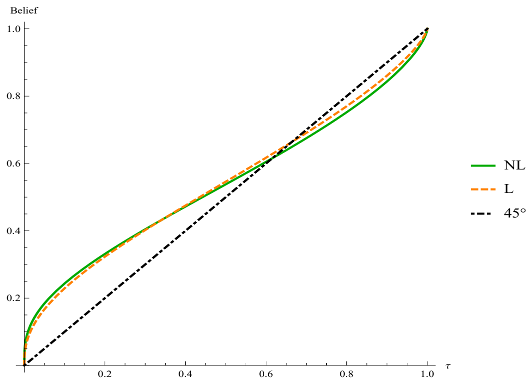

Figure 2.4: Estimated probability weighting function for unambiguous probabilities in NL & L.

ous priors. Assuming that the pattern does not change when information is released gradually, this pattern is used to discuss the validity of FBU and MLU posteriors.

The probability weighting function has single probability values as an argument. Ambiguous distributions, however, are characterized by intervals of probabilities. I approximate the estimates of the weighting function by using the midpoint of the set of probabilities. In treatment NL, the set corresponds to the ambiguous set of priors $[\pi_l, \pi_h]$. In treatment L, the set equals the set of posteriors that varies with the updating rule. The midpoints of the set of FBU posteriors are less extreme than the midpoints of the set of MLU posteriors, which here is a singleton.

The solid line in Figures 2.5a and 2.5b depicts the relation between subjects' estimated beliefs and ambiguous probabilities in NL. This inverse s-shape relation serves as benchmark for the relation between estimated beliefs and ambiguous posterior probabilities in L. The dashed line in Figure 2.5a represents the model fit with FBU posteriors. Estimated beliefs are slightly s-shaped in FBU posteriors, rather than inverse s-shaped. The discrepancy between the benchmark (solid line) and the fit with FBU posteriors (dashed line) points out that decision weights with FBU posteriors are too extreme. That is, trading decisions are too extreme to be explained by the range of beliefs under FBU.

The dashed line in Figure 2.5b depicts the model fit with MLU posteriors. The weighting function is inverse s-shaped, but deviates from the benchmark (solid line) as well. Given MLU probabilities, estimated beliefs are not sufficiently extreme to match the benchmark. Trading decisions are too close to the belief of 50% to be explained by extreme MLU posteriors. Section B.3.2 in the Appendix displays the estimates of the NNLS-SUR with ambiguous probabilities and the results of a Lagrange-Multiplier test, which shows a significant difference between the benchmark model and the model fit under both FBU and MLU probabilities.

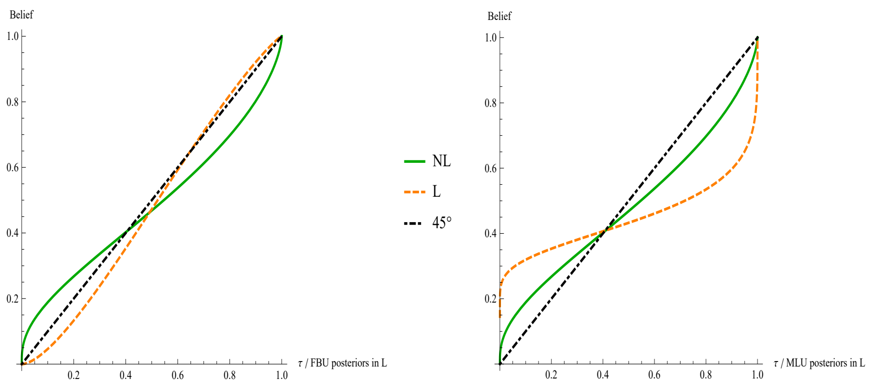

(a) Assuming FBU in L. **(b)** Assuming MLU in L.

Figure 2.5: Estimated probability weighting function for ambiguous probabilities in NL & L.

In a nutshell, quotes based on ambiguous posteriors are not extreme enough to be explained by MLU beliefs, but too extreme to be explained by FBU beliefs.

Heterogeneous updating rules

Heterogeneity in updating behavior partly accounts for the bad fit of FBU and MLU models. To illustrate the heterogeneity in quotes, the midpoints of bid and ask pairs (henceforth mid-quotes) are depicted in Figure 2.6. The top two panels 2.6a & 2.6b show the distribution of mid-quotes for the ambiguous probabilities $\pi \in [.05, .65]$ and $\pi \in [.35; .95]$, respectively. Without incoming information, mid-quotes are distributed symmetrically around the midpoint of the set of probabilities. The distributions differ clearly in the bottom two panels 2.6c & 2.6d, that show mid-quotes for the same intervals of FBU posteriors (i.e. $\rho \in [.05, .65]$ and $\rho \in [.35; .95]$). Mid-quotes are clustered at 3 mass points ($\{0-5; 20-25; 45-50\}, \{50-55; 70-75; 95-100\}$) suggesting 3 main updating methods.

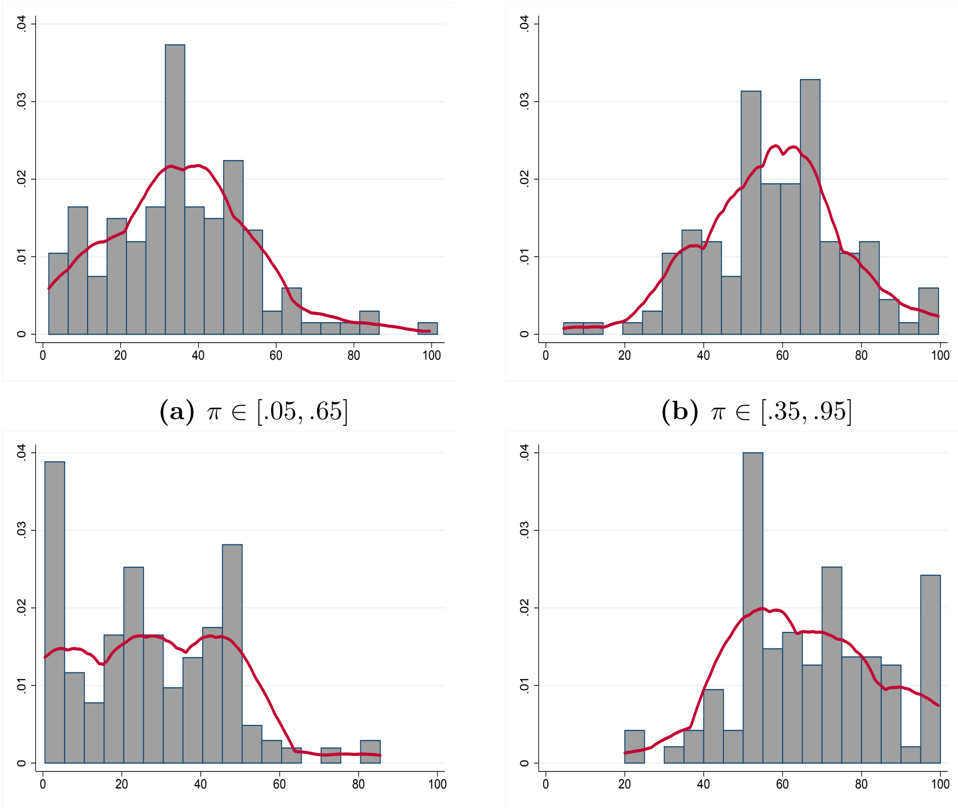

(a) $\pi \in [.05, .65]$

(b) $\pi \in [.35, .95]$

(c) $\rho \in [.05, .65]$

(d) $\rho \in [.35, .95]$

Figure 2.6: Mid-quotes for π or $\rho \in [.05, .65]$ (left) and π or $\rho \in [.35, .95]$ (right). Treatment NL in top panels, L in bottom panels.

The cluster analysis in Appendix B.3.3 illustrates how trading decisions differ. In sum, a substantial share of quotes (25.75%) match highly ambiguity-averse trading behavior, that favors non-participation. These subjects center their bids and asks around the mid-prior 50 and choose wide spreads. Another substantial share (21.72%) is consistent with MLU: They choose extreme quotes and minimal spreads. The majority of trading decisions (42.17%) is consistent with less extreme Bayesian quotes. However, these Bayesian quotes do not reflect FBU posteriors. Under FBU, participants in treatment NL and L should consider

the same support of probabilities and therefore make similar trading decisions. Bid-ask pairs in treatment L should resemble the ones in NL and should be similarly centered around the midpoints of the sets of probabilities, which are here $\{35, 65\}$. However, the bid-ask quotes based on incoming information encompass beliefs that are more extreme than the ones in treatment NL. Controlling for the range in marginal and FBU probabilities, 36.19% of bid-ask pairs in treatment NL encompass the value 50 versus 29.54% in treatment L (p-value=0.07 in binomial test). Table B6 in Appendix B.3.3 shows that the same cluster analysis in treatment NL provides ranges of bid-ask pairs that are less extreme, with observations that are distributed more evenly across the clusters.
Chosen quotes can be rationalized with the updating of the prior $\pi = .5$, the midpoint of the set of priors. Indeed, the Bayesian posterior $\rho(s, \pi = .5)$ fits the relation between trading decisions and probabilities (see Figure 2.7): The probability weighting functions with marginal and posterior probabilities do not differ, when posterior probabilities correspond to Bayesian updates of the midpoint of ambiguous priors.[9]

In a nutshell, generally subjects are ambiguity-averse, even after receiving information about ambiguous priors. Learning does not impact the spread, but induces different and more heterogeneous quotes. Furthermore, Bayesian updates of the mid-prior describe aggregate quotes better than FBU or MLU posteriors. These results point out the relevance of conditional smooth preferences. The next section outlines to what extent Bayesian smooth preferences explain data.

2.4.3 Conditional smooth preferences

This section shows that chosen quotes in the two treatments are consistent with second-order preferences over probabilities. In a first step, I show that ambiguity-averse, but midpoint-preserving recursive preferences generate a bid-

[9] Since the conditional probability for a correct signal is $q = .75$, the mass points around 25 and 75 suggest base-rate neglect as a possible explanation. However, base-rate neglect is unlikely to cause this pattern. Base-rate neglect should become apparent in decisions regarding both ambiguous and unambiguous return distributions. Yet, subjects - even those who fall in this specific cluster of Bayesian updaters - adjust their quotes to the prior in risky rounds. Figure B6 in Appendix B.3.2 shows how mid-quotes increase in the prior for the different signal values. Bids and asks are not heavily centered around 25 or 75.

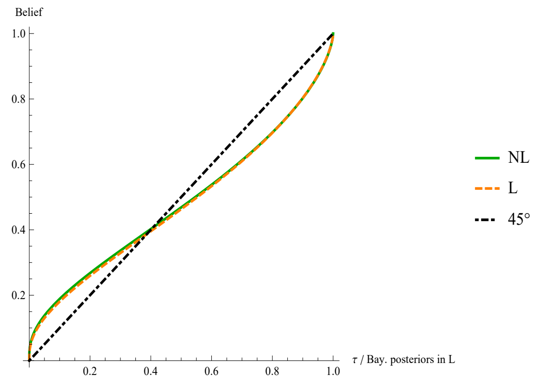

Figure 2.7: Estimated probability weighting function for ambiguous probabilities in NL & L assuming BU of mid-prior in L.

ask spread around the midpoint of possible priors. In a second step, I show that quotes match Bayesian updates of recursive preferences.

Following the model of smooth preferences in Klibanoff et al. (2005), a strictly increasing and concave function $\phi(\cdot)$ is used to represent ambiguity-averse second-order preferences. The agent's value function is assumed to take the double expectational form:

$$\int_{\pi_l}^{\pi_h} \phi(\mathbb{E}_\pi U(\cdot))\psi(\pi)d\pi \tag{2.3}$$

where $\psi(\pi)$ represents the subjective probability over the set of priors $[\pi_l; \pi_h]$. The operator $\mathbb{E}_\pi$ computes the expected value with respect to a specific Bernoulli distribution $f(\pi)$ with success probability π.

Like in standard expected utility models, attitudes towards risk are captured by the concavity of a von Neumann-Morgenstern utility function $U(\cdot)$. In addition, attitudes towards ambiguity are captured separately by the function $\phi(\cdot)$. Agents assign subjective second-order beliefs $\psi(\pi)$ to some probability distribution π. In their decision-making, they evaluate subjective expectations over expected utilities. Ambiguity aversion corresponds to a dislike of spreads around the mean expected utility and is reflected by the concavity of the function $\phi(\cdot)$. The following analysis assumes that subjects have second-order beliefs, whose mean corresponds to the midpoint in the range of priors. This assumption is in line with the principle of insufficient reasons, under which agents assign equal probabilities to mutually exclusive events if they have no explicit reason to do differently.[10]

I first show that ambiguity-averse second-order preferences generate a bid-ask spread. The optimal bid for going long is the certainty equivalent that satisfies:

$$\int_{\pi_l}^{\pi_h} \phi(\mathbb{E}_\pi U(W_0 + V - b))\psi(\pi)d\pi = \phi(U(W_0)) \tag{2.4}$$

Denote $\int_{\pi_l}^{\pi_h}(\cdot)\psi(\pi)d\pi =: \mathbb{E}_\psi(\cdot)$. By Jensen's inequality:

$$\mathbb{E}_\psi \phi(\mathbb{E}_\pi U(W_0 + V - b)) < \phi(\mathbb{E}_\psi \mathbb{E}_\pi U(W_0 + V - b)) \tag{2.5}$$

Under mean-preserving second-order beliefs, the subjective probability functions $\psi(\pi)$ satisfies $\int_{\pi_l}^{\pi_h} \pi\psi(\pi)d\pi = \mathbb{E}_\psi(\pi) = \bar{\pi}$, where $\bar{\pi}$ represents the midpoint of priors. The RHS in Equation (2.5) equals then:

$$\phi(\mathbb{E}_\psi \mathbb{E}_\pi U(W_0 + V - b)) = \phi(\mathbb{E}_{\bar{\pi}} U(W_0 + V - b)) \tag{2.6}$$

Consider an agent who bids for a risky asset with Bernoulli distribution $f(\bar{\pi})$. The optimal bid makes the agent indifferent between buying the asset and keeping the endowment. It satisfies :

$$\phi(\mathbb{E}_{\bar{\pi}} U(W_0 + V - b^R) = \phi(U(W_0)). \tag{2.7}$$

From equations (2.4), (2.5) and (2.7) it follows that:

$$\phi(\mathbb{E}_{\bar{\pi}} U(W_0 + V - b^R) < \phi(\mathbb{E}_{\bar{\pi}} U(W_0 + V - b)). \tag{2.8}$$

Because $\phi(\cdot)$ is strictly increasing, $U(\cdot)$ strictly concave, the optimal bid under ambiguity aversion is smaller than the optimal bid under risk, $b^{AA} < b^R$. Analogously, $a^{AA} > a^R$. Ambiguity-averse smooth preferences produce wider spreads

[10]Henceforth, the notion "mean-preserving" refers to "midpoint-preserving" in this context.

than the spread under risk. With mean-preserving second-order beliefs, bid and ask quotes converge to the expected value under risk with decreasing ambiguity and risk aversion.

Incoming information alters the optimization problem at two points. First, expected utility is computed with posterior probabilities $\rho(s,\mu)$ instead of given probabilities π. Second, the incoming information affects directly second-order beliefs $\psi(s,\mu)$ by shifting more weight to more likely probability values (Epstein and Schneider, 2007; Klibanoff et al., 2009). With standard Bayesian updating:

$$\psi(s,\mu) = \frac{\psi(\mu)f(s,\mu)}{\int_{\mu_l}^{\mu_h} \psi(\tilde{\mu})f(s,\tilde{\mu})d\tilde{\mu}}$$

where

$$f(s,\mu) = \begin{cases} q\mu + (1-q)(1-\mu) & if\ s = \vartheta_h \\ (1-q)\mu + q(1-\mu) & if\ s = \vartheta_l \end{cases}$$

The function $f(s,\mu)$ is the probability of receiving signal s given a prior Bernoulli distribution with success probability μ. In particular, because $\psi(s,\mu) \neq \psi(\mu)$:

$$\mathbb{E}_{\psi(s=\vartheta_l,\mu)}\phi(\mathbb{E}_{\{s=\vartheta_l,\mu\}}U(\cdot)) < \mathbb{E}_{\psi(\mu)}\phi(\mathbb{E}_{\{s=\vartheta_l,\mu\}}U(\cdot)) \tag{2.9}$$

$$\mathbb{E}_{\psi(s=\vartheta_h,\mu)}\phi(\mathbb{E}_{\{s=\vartheta_h,\mu\}}U(\cdot)) > \mathbb{E}_{\psi(\mu)}\phi(\mathbb{E}_{\{s=\vartheta_h,\mu\}}U(\cdot)) \tag{2.10}$$

Therefore, $b^{CSP,}_{\{s=\vartheta_l\}} < b^{SP}_{\{s=\vartheta_l\}}$: With conditional smooth preferences (CSP), second-order beliefs over priors that are updated upon the signal ($s = \vartheta_l$) induce a bid b^{CSP} that is lower than the optimal bid obtained with the same second-order beliefs over marginal probabilities. Analogously, $b^{CSP,}_{\{s=\vartheta_h\}} > b^{SP}_{\{s=\vartheta_h\}}$. Thus, conditional smooth preferences generate more extreme beliefs than marginal smooth preferences if traders have mean-preserving second-order beliefs. Consequently, gradual information release induces more extreme quotes compared to an environment where information is released all at once. Figures 2.8a and 2.8b display second-order beliefs with and without learning for the same support of probabilities. The dashed line depicts a uniform density over probabilities, which can be interpreted as subjects' uniform second-order beliefs over marginal probabilities (applicable to treatment NL). The solid lines represent second-order beliefs over posteriors after Bayesian updating of uniform second-order beliefs over priors (applicable to treatment L). With smooth preferences, final expectations are more extreme if information is learned progressively.

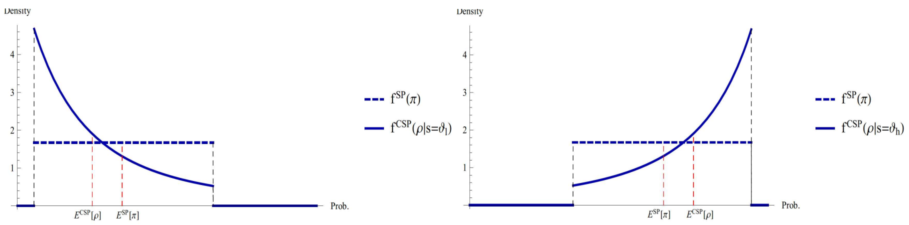

(a) $\pi, \rho \in [.05, .65]$ **(b)** $\pi, \rho \in [.35, .95]$

Figure 2.8: Marginal and Bayesian second-order beliefs for a low (a) and a high (b) support of probabilities.

In addition, it can be shown that under the assumption of mean-preserving spreads: $b^{CSP} < b^R$. The risk-neutral agent quotes: $b^{RN} = a^{RN} = E(V|s, \bar{\mu})$, where $\bar{\mu} = E[\mu]$. With decreasing ambiguity and risk aversion: $b^{CSP} \longrightarrow E(V|s, \mu = E[\mu])$. Analogously, $a^{CSP} > a^R$ and $a^{CSP} \longrightarrow E(V|s, \mu = E[\mu])$ with decreasing ambiguity and risk aversion.

With the principle of insufficient reasons, for instance, the mean prior belief corresponds to $E[\mu] = .5$ for $\mu \in [.15, .85]$. Bids and asks would be centered around $E[V|s, \mu = .5]$, i.e. $E[V|s = \vartheta_l, \mu = .5] = 25$ after a low signal and $E[V|s = \vartheta_h, \mu = .5] = 75$ after a high signal. In this context, Bayesian updating of second-order preferences can explain why quotes are more extreme in treatment L than in NL for the same support of probabilities.

2.5 Conclusion

The evidence of ambiguity aversion found so far in Ellsberg-type experiments extends to other frameworks. The experiment shows that, in cases where portfolio reallocation is limited, ambiguity impedes willingness to trade - with and without sequential information processing. These results confirm the intuition that investors appear to consider ambiguous assets as more risky (Sarin and Weber, 1993; Epstein and Wang, 1994).

A second main insight of the experiment is that ambiguity effects cannot be disentangled from the information condition. The same degree of ambiguity leads to different trading decisions, depending on how many pieces of information have been available so far. Despite the same willingness to trade, investors choose more extreme quotes when they receive information in pieces.

In addition, incoming information introduces more heterogeneity in trading behavior. A substantial share of agents is insensitive to additional information, another non-negligible share adopts extreme beliefs, and the majority of agents appears to update second-order preferences in a Bayesian way.

The heterogeneity in information processing advises caution on the general conclusion in Baillon et al. (2013) "[...] that pessimism is a stable trait of decision makers, not affected much by information received." Indeed, the aggregate distribution of spreads remains stable when information is released sequentially. However, the heterogeneity in updating rules shows that the stability of ambiguity attitudes may not hold for everyone. This is in line with Bossaerts et al. (2010), who argue that heterogeneity has important implications for markets, which are, therefore, not best described by a representative agent. Heterogeneity in updating behavior, though, probably impacts markets differently than heterogeneity in ambiguity attitudes. For instance, heterogeneity in trades might be amplified if specific traders are more prone to use specific updating rules. This raises the question of whether the updating rule is inherently influenced by ambiguity preferences. If highly ambiguity-averse subjects are insensitive to information and less ambiguity-averse subjects are instead more prone to apply MLU, gradual information flow may reinforce disparities in ambiguity preferences. In particular, asset pricing would be determined by extreme updaters, if those who update cautiously refrain from trading.

Other important questions remain to be clarified in future research. First, ambiguity effects possibly differ in markets. There is a difference between individual willingness to trade and its counterpart in markets, e.g. liquidity or market depth. The risk of adverse selection may incite investors to avoid ambiguous markets even more. Alternatively, trade may possibly be driven by one's knowledge relative to other market participants (Zeckhauser 2006, cf. competence hypothesis in Heath and Tversky 1991). To be willing to trade, it might be sufficient to be not at informational disadvantage compared to other traders. Furthermore, the interaction between investors might eliminate any perception of ambiguity, especially if markets are dominated by aggressive traders. The findings in Sarin and Weber (1993), though, indicate that ambiguity effects are robust to market feedback. Yet, the extent to which information aggregation abates ambiguity effects is still not clear.

Second, the observed divergence in beliefs casts doubts on the hypothesis that trading volume falls with ambiguity. Even if ambiguity weakens individual willingness to trade, beliefs resulting from learning might be so divergent that different trading parties agree on speculative trade.

Third, updating behavior may vary with the type of information. Information in itself can be ambiguous. Extreme updating possibly disappears when the precision of signals is not known. The reactions to information matters in particular, if the decision to acquire information is endogenous. A correlation between willingness to pay for information under ambiguity and a subject's ambiguity preferences might abate or reinforce ambiguity effects.

In sum, it is important to identify conditions under which ambiguity effects are self-enforcing. A faster resolution of ambiguity and a concomitant increase in liquidity benefit not only trading venues through higher profits, but also investors through lower transaction costs and, potentially, higher price efficiency. This study draws the attention to frequent information release as a mechanism to avoid or correct frictions in trades.

3 Preference-dependent learning

3.1 Introduction

Decisions under risk are characterized by heterogeneity. Heterogeneity in decisions might be driven by random errors in decision-making, by variation in preferences, or variation in beliefs. There is robust evidence of heterogeneity in preferences, in particular with respect to risk (Bruhin et al., 2010; Conte et al., 2011). In addition, agents hold different beliefs because of information asymmetry. But even in markets with symmetric information, beliefs may differ once agents process incoming information. Differences in updating behavior are documented, for instance, in Grether (1980, 1992), and El-Gamal and Grether (1995).

While random errors are likely to cancel out in aggregate outcomes, heterogeneous preferences or beliefs are shown to impact equilibria in some markets. Chapman and Polkovnichenko (2009), for instance, demonstrate how heterogeneous preferences affect risk premia and risk-free rate when agents depart from expected utility. Haltiwanger and Waldman (1985, 1989) show that, under strategic complementarities, heterogeneous abilities to process information alter equilibrium outcomes. More generally, market equilibria are often derived with representative agents. The degree of heterogeneity, though, might determine whether and how quickly markets converge to these defined equilibria.

Rather than focusing on the consequences, the present paper examines the origins of heterogeneity. To this end, I estimate the importance of preferences in decision-making and study the interaction between preferences and belief formation. Depending to what extent belief formation correlates with preferences, markets that require belief revision might exhibit more or less variance in decisions.

To the best of my knowledge, no previous study addresses the question of a systematic link between preferences and belief updating. One of the reasons is that expected utility theory (EUT), the main paradigm in economic theory, assumes separability of utilities and probabilities. Yet, a large body of experimental evidence proves this assumption wrong.[1] Allais (1953), for instance, points to the common ratio effect, in which a downward scaling in probabilities often reverses choices. Moreover, the experimental literature on risk attitudes robustly documents this fourfold-pattern: Subjects are, on average, risk-seeking for small (large) but risk-averse for large (small) probabilities in the gain (loss) domain. Such systematic violations of EUT have prompted the development of several non-expected utility theories that model an alternative weighting of the valuation of outcomes (e.g. disappointment aversion, Bell 1985; Loomes and Sugden 1986), probabilities (e.g. subjective expected utility (SEU), see Fishburn 1981 and references therein), rank-dependent utility (RDU, Quiggin 1982), or both (e.g. prospect theory, Kahneman and Tversky 1979).

Using the experiment described in Chapter two (cf. Ngangoue 2016), I find supportive evidence of an interaction between risk preferences and belief revision. To arrive at this conclusion, I first estimate the correlation between risk preferences and investment decisions in the case where decision-makers (DMs) know the final distribution of states of nature and then investigate whether the effect is strengthened or dampened when belief updating comes into play.

An ubiquitous problem with elicited preferences is that their variance is driven to a substantial part by measurement error (e.g. 30-40% in the data analysis of Gillen et al. 2015). After correcting for measurement error with simulation extrapolation (Cook and Stefanski, 1994; Devanarayan and Stefanski, 2002), the weight of risk attitudes in investment decisions not only deviates from the theoretical benchmark, but also differs across the information condition. Without learning, an increase in risk aversion induces a moderate increase in spreads. Upon belief revision, in contrast, spreads react more to increased aversion. The substantial change in estimates casts doubts on the independence assumption between risk preferences and updating behavior.

Motivated by this observation, I test whether belief revision deviates from Bayes rule for different types of risk attitudes. While decisions of risk-seeking and risk-neutral subjects conform with Bayesian inference, risk-averse subjects display

[1] See Kahneman and Tversky (1979); Loomes and Sugden (1983, 1987); Starmer and Sugden (1989); Battalio et al. (1990); Conlisk (1989), inter alia.

more conservatism in their decisions. In this experiment, differences in preferences determine the variance in decisions, not only directly but also indirectly through belief revision. Further inquiries are needed, but neglecting this interaction may result in underestimating the importance of risk preferences for decision-making with learning dynamics.
The following section describes the data, which consists of investment decisions and elicited risk preferences. Section 3.3 defines the decision model and the hypothesis to be tested. Section 3.4 presents results, followed by a test on independence between learning abilities and risk preferences in Section 3.5. Section 3.6 concludes.

3.2 Data

The analysis is conducted with the individual decision-making experiment described in Chapter two (Ngangoue, 2016). The computerized experiment consisted of two treatments, called "No Learning" and "Learning," each containing two parts. In Part 1, subjects made several investment decisions. After all subjects had completed Part 1, their risk preferences (among other measures) were elicited in Part 2. Treatment "No Learning" had 67 participants, while a different pool of 66 subjects participated in treatment "Learning." The difference between the two treatments is explained in the following description of Part 1.

3.2.1 Investment decisions

In Part 1, subjects have the possibility to invest in a risky asset. The asset value, $\theta \in \{0,1\}$, follows a Bernoulli distribution with success probability π, $B(\pi)$.[2] Before making any decision, subjects receive information about the Bernoulli distribution. In the treatment "No Learning" (NL), they learn the final success probability $\pi = Pr(\theta = 1)$. In the treatment "Learning" (L), subjects first learn a prior probability μ and then receive an additional binary signal $s \in \{0,1\}$. The signal's precision is symmetric and given by the inverse

[2]For simplicity, the experimental design and the following analysis are described with normalized values between 0 and 1. The actual experiment was conducted with parameters between 0 and 100, which, in turn, make the current analysis cumbersome without providing any further insights.

of $q = Pr(s = \theta|\theta) = 0.75$.
Subjects in both treatments submit a bid and an ask quote, $\{b, a\} \in [0, 1]$. A random price $p \in [0, 1]$ determines subsequently whether a buy, a sell, or no trade occurs. A subject buys one share if the asset price falls below her bid, but short-sells it if the price exceeds her ask. No trade occurs if a subject states a spread between bid and ask, within which the price realizes. The variable X summarizes subject's demand:

$$X = \begin{cases} +1 & \text{if } p \leq b \\ 0 & \text{if } b < p < a \\ -1 & \text{if } p \geq a \end{cases} \tag{3.1}$$

Under expected utility maximization, risk-neutral DMs set $b = a = E(\theta)$ (cf. Ngangoue 2016). Quotes of risk-averse DMs, however, include a risk premium $\Gamma(\pi)$, leading to $b = E(\theta) - \Gamma(\pi) < E(\theta)$ and $a = E(\theta) + \Gamma(1 - \pi) > E(\theta)$.
In each treatment, the investment task is repeated across 14 independent rounds with varying Bernoulli probabilities π or $\mu \in [0.05, 0.95]$.[3] Subjects begin each round with a cash endowment W_0 of 1 currency unit. Figures C1a and C1b in Appendix C.1 show the distribution of chosen spreads in each treatment.

3.2.2 Risk preferences

Risk preferences are elicited with a multiple price list task akin to Abdellaoui et al. (2011) and Gillen et al. (2015). In two replicate measurements, subjects face a list of pairwise choices between a sure payoff and a lottery. Define the lottery $(x, \pi; 0)$ as the chance to win prize x with probability π, and win nothing else. The lotteries in the first and second measurement correspond to $(100, 0.5; 0)$ and $(150, 0.5; 0)$, respectively. The lottery is illustrated on the left side of the computer interface, where subjects see an urn with 10 (15) yellow and 10 (15) black balls in the first (second) measurement. The lottery pays out if a black ball is drawn. The right side of the interface shows a list of sure payoffs in $[0; x]$, with increments of 5 ECU per row. Subjects must then, for each row, make a

[3]The exact probabilities are as follows: $\{0.05; 0.15; 0.35; 0.5; 0.65; 0.85; 0.95\}$. Every probability value is used twice. This procedure leads to posterior probabilities between $[0.02; 0.98]$ in treatment L (see Table 2.1 in Chapter 2).

pairwise choice between the lottery and the sure payoff. Monotonicity is enforced as subjects can only switch once from preferring the lottery to preferring a sure payoff. Figure C2 in Appendix C.1 depicts the computer interface for the first measurement with lottery $(100, 0.5; 0)$. Feedback on payoff was provided only after completion of Part 2.

The certainty equivalent (CE) is the lowest sure payoff that the DM prefers to the lottery. Figure C3 in Appendix C.1 depicts the distribution of relative risk premia ($RPP = \frac{E(x)-CE}{E(x)}$), averaged across both measurements. On average, 57.9% of all subjects choose a positive risk premium, exhibiting thereby risk aversion.

The decision-model described in Section 3.3 assumes constant relative risk aversion (CRRA). For consistency, imputed CRRA coefficients γ serve as a measure of risk aversion and are obtained by solving the equation $U(CE) = EU(Lottery)$ with $U(z) = \frac{z^{1-\gamma}}{1-\gamma}$ (see summary statistics in Table 3.1). The relevant range of CRRA coefficients is given by the CE in two rows: the last row for which the DM prefers the lottery and the first row for which she prefers the sure payoff. For instance, if she plays the lottery up to a sure payoff of 40 ECU and then prefers a sure payoff of at least 45 ECU, the relevant range of her risk attitude is given by imputed coefficients with a CE of 40 and 45 ECU. Following standard procedures, the analysis is then conducted with the arithmetic mean of the two imputed CRRA coeffcients (e.g. Habib et al. 2017).

Dealing with inconsistent responses. Subjects who always prefer the lottery or the sure payoff are excluded from the analysis. Choosing the safe or the risky option for all rows could be rationalized by pure indifference or erroneous reasoning, but is unlikely to be an indicator of extreme preferences. This procedure leads to the exclusion of 3 subjects in treatment L. Furthermore, I assume that single extreme choices of no switching - more present in the first measure - stem from mistakes or curiosity. To limit their bias on the arithmetic mean, I replace 6 and 2 of these observations in NL and L by their duplicates.

Subjects in treatment L display less risk-aversion than the ones in treatment NL (see Table 3.1). At this point, however, it is not possible to discern whether

Table 3.1: Summary statistics of imputed CRRA coefficients

	NL & L	NL	L
Median γ	0.0556	0.1296	-0.0172
% averse	55.38	62.69	47.62
Mean $\mid \gamma > 0$	0.3225	0.3155	0.3325
Median $\mid \gamma > 0$	0.3171	0.3014	0.3380
St. Dev.	0.21	0.19	0.24
N	130	67	63

Note: Inconsistent responses are excluded. Statistics correspond to mean values of the two replicate measures. Medians between the treatments differ at marginal significance (p=0.05 for overall medians in non-parametric median test), but do not differ when conditioning on positive values (p=0.811 in median test).

this difference is due to random difference in sampling or whether preference elicitation is directly affected by Part 1 of the experiment. Conditional on risk aversion (i.e. $\gamma > 0$) though, treatments do not significantly differ (p=0.811 in non-parametric median test).

Eliciting preferences in two replicate tasks has one main advantage. It enables to assess how much of the variance in responses is driven by random errors. For a crude assessment of the error size, assume that the individual coefficient γ_i follows a distribution with mean $E[\gamma_i] = \gamma$ and variance $Var[\gamma_i] = \sigma_\gamma^2$. Furthermore, assume that the replicate measures are error-contaminated variables of the true risk aversion parameter γ_i: $\tilde{\gamma}_{ij} = \gamma_i + \sigma_i u_{ij}$, with $E[\tilde{\gamma}_{ij}] = \gamma_i$, $Var[\tilde{\gamma}_{ij}] = \sigma_\gamma^2 + \sigma_i^2$, $E[u_{ij}] = 0$ and $E[u_{il}u_{ik}] = 0$ for $l \neq k$. The correlation between replicate responses can be used to estimate the importance of measurement error σ_u in elicited preferences. In the case of homoscedasticity ($\sigma_i^2 = \sigma_u^2$, $\forall i$), the correlation coefficient is given by $\tau = Corr(\tilde{\gamma}_{i1}, \tilde{\gamma}_{i2}) = \frac{\sigma_\gamma^2}{\sigma_\gamma^2 + \sigma_u^2}$. Rearranging the equation provides an estimate for the proportion of the error variance relative to the variance in preferences: $\sigma_u^2 = \frac{1-\tau}{\tau}\sigma_\gamma^2$. In the experiment, the low correlation between the replicate measures ($\tau = 0.66$) discloses substantial choice switching across tasks.[4][5] In other words, the measurement error σ_u^2 accounts for a third of the variance in elicited preferences ($\sigma_u^2 \approx \frac{1}{2}\sigma_\gamma^2$, implying $Var[\tilde{\gamma}_{ij}] \approx \frac{3}{2}\sigma_\gamma^2 = 3\sigma_u^2$).[6] This within-subject variance highlights the necessity to correct for measurement error.

It needs to be emphasized that the two risk measures are an indication for risk preferences, but have serious limitations (Fehr-Duda and Epper, 2012). Due to time constraints and an interest of keeping instructions as simple as possible, risk preferences were elicited using prospects with single non-zero outcomes and a fixed probability of 50%. Consequently, nothing can be inferred about individual, non-linear probability weights. The individual utility curvature is also not assessable, unless a functional form for utils is imposed. Nevertheless, these limitations hold for both treatments and are unlikely to account for major differences between them.

Payoffs were determined with a random incentive system at the beginning of Parts 1 and 2 (Baillon et al., 2015). That is, at the beginning of each part, a round or task was randomly drawn for payoff. The exchange rate was 0.13 EUR per ECU.[7] Subjects' earnings were between 5 and 28.84 EUR, with an average of 19.50 EUR for approximately 100 minutes.

[4]Pearson's correlation coefficients τ in the total sample, treatment NL, and treatment L correspond to 0.66, 0.502, and 0.84, respectively.

[5]Similar fluctuations in choices are observed in various experimental studies (Mosteller and Nogee, 1951; Hey, 2001; Camerer, 1989; Hey and Orme, 1994; Starmer and Sugden, 1989; Ballinger and Wilcox, 1997; Loomes and Sugden, 1998). Choices are also found to be sensitive to the elicitation procedure (Andersen et al., 2006; Lévy-Garboua et al., 2012). Some errors in the present experiment are likely to come from the fact that preferences were elicited after the investment task, where subjects' attention might not have been as high anymore.

[6]The relative sizes of the error variance, $\frac{1-\tau}{\tau}$, in the total sample, treatment NL, and treatment L equal 0.53, 0.99, and 0.19, respectively.

[7]The exchange rate is defined for the original parameters of the experiment, which, here, have been normalized to $[0, 1]$ in the description of Part 1.

3.3 Decision model under EUT

3.3.1 Theoretical model

The analysis is based upon a decision model that postulates a relation between chosen spreads in the investment task and the risk aversion measure γ. Using the Arrow-Pratt approximation, the optimal bid and ask is approximated by:

$$b^* = \pi - \frac{1}{2}\, A(W_0) \cdot V_\theta \tag{3.2}$$

$$a^* = \pi + \frac{1}{2}\, A(W_0) \cdot V_\theta \tag{3.3}$$

where $A(W_0)$ denotes the Arrow-Pratt measure of absolute risk aversion at initial wealth W_0 (the derivation can be found in Appendix C.2).[8] The term $V_\theta = \pi(1 - \pi)$ corresponds to the variance of the asset value when the DM holds belief π. For the class of utility functions with coefficient of constant relative risk aversion γ, the optimal spread is approximated (for small γ) by:

$$a^* - b^* = \max(0, m(\gamma, V_\theta)) \tag{3.4}$$

$$\texttt{with}\ \ m(\gamma, V_\theta) = \gamma\, V_\theta$$

The optimal spread is the investment risk, V_θ, weighted by the DM's risk attitude γ. Equation (3.4) illustrates in particular the separability between probability and risk preferences under EUT: The coefficient of CRRA, γ, captures the curvature in the utility function, while the risk of the investment, V_θ, only depends on probability beliefs.

3.3.2 Econometric model

Decision model (3.4) provides the structural model for estimating the weight of the two relevant factors in subjects' decisions: risk preferences and investment risk. To this end, a non-linear two-limit Tobit regression is used to estimate the

[8] The optimal bid and ask do not lie exactly symmetric around belief π, but for $|\gamma| \to 0$, the Arrow-Pratt approximation converges toward the optimal quotes. Here, for the majority of risk-averse subjects with $\gamma \leq 0.5$, the Arrow-Pratt approximation deviates from the optimal spread by less than 0.0001 ECU.

following latent variable model for the observed spread y_i:

$$y_i = a_i - b_i = \begin{cases} 1 & if\ {y_i}^* \geq 1 \\ {y_i}^* & if\ \underline{y_i}^* < {y_i}^* < 1 \\ 0 & if\ {y_i}^* \leq \underline{y_i}^* \end{cases} \tag{3.5}$$

where $y_i^* = e^{m(\gamma,V_\theta)+\epsilon_i} = e^{\beta_0+\gamma_i^{\beta_1}\,{V_\theta}^{\beta_2}+\epsilon_i}$, $\underline{y_i}^* = e^{m(0,V_\theta)}$ and $\epsilon_i \sim N(0,\sigma_\epsilon)$.

The latent decision variable y_i^* is a function of $m(\gamma, V_\theta)$ and some disturbance ϵ_i. In the econometric analysis, the function $m(\cdot)$ takes the general form $m(\gamma, V_\theta) = \beta_0 + \gamma^{\beta_1}\,{V_\theta}^{\beta_2}$, where the coefficient vector $(\beta_0, \beta_1, \beta_2)'$ contains the main estimates of interest. For $(\beta_0, \beta_1, \beta_2)' = (0, 1, 1)'$ the function $m(\gamma, V_\theta)$ coincides with decision model (3.4), but DMs may also assign alternative weights to their inherent risk preference and the objective investment risk.

The spread, which has a right-skewed distribution, is assumed to be lognormal. The variable y_i^* takes therefore the exponential form. The DM chooses a maximum spread of 1 when y_i^* exceeds the value 1. Analogously, she chooses zero spread when y_i^* falls below a lower threshold $\underline{y_i}^*$. Theoretically, only risk-averse DMs, i.e. DMs with $\gamma_i > 0$, choose a positive spread whereas $y_i = 0$ for all DMs with $\gamma \leq 0$. Thus, with strictly positive investment risk ($V_\theta > 0$), the lower threshold $\underline{y_i}^*$ is given by $e^{m(0,V_\theta)}$.

The two error-contaminated CRRA coefficients $\tilde{\gamma}_{ij},\ j = \{1, 2\}$, serve as a measure of risk attitude and are integrated in the estimation with the following assumptions: $\tilde{\gamma}_{ij} = \gamma_i + \sigma_i u_{ij}$ where $u_{ij} \sim i.i.d\ N(0, 1)$.[9][10]

Measurement error is extrapolated with the simulation extrapolation (SIMEX) procedure for heteroscedastic errors with unknown variances and replicate measurements (Devanarayan and Stefanski, 2002; Carroll et al., 2006). The SIMEX

[9] Because the function $m(\cdot)$ is not continuous in β_1 for $\gamma < 0$, I estimate the Tobit model for $(\frac{\gamma+\Delta}{\Delta+1})$. Consequently, I interpret marginal effects rather than coefficients $(\beta_0, \beta_1, \beta_2)$ directly.

[10] The model is estimated assuming $((\sigma_i u_{ij})\epsilon_i)' \sim N(\mathbf{0}, \mathbf{\Sigma}_i)$ with $\mathbf{\Sigma}_i = \begin{pmatrix} \sigma_i^2 & 0 \\ 0 & \sigma_\epsilon^2 \end{pmatrix}$. The assumption that disturbances $(\sigma_i u_i, \epsilon_i)$ in the investment task and the preference elicitation task are uncorrelated might be strong, but there is no indication of a significant correlation in the data.

procedure is applicable to various estimation methods and is approximately consistent in nonlinear models. Its basic idea is to estimate the relation between error variance and parameter estimates. To this end, different data sets are created by inflating to various extents the error-contaminated variables with additional error.[11] Estimates are then obtained for each of these variance-inflated data sets, and allow for estimating the relation between error variance and corresponding estimates. By means of a polynomial extrapolation, this link function is finally extrapolated to the case where error variance equals zero (see Appendix C.3 for more details).

3.3.3 Hypothesis

Treatment NL, where no learning occurs, serves as benchmark for the general weights of risk attitudes in model (3.5). In case of independence between risk preferences and belief updating, coefficients for risk preferences in treatments NL and L should be equal.

Hypothesis 4 *Under EUT:*

$$H_0 : \quad \beta_1^{NL} = \beta_1^{L}$$

where the coefficient β_1 renders the weighting of risk preferences.

Any interaction between preferences and updating behavior, in contrast, would bias β_1. With preference-dependent learning, the variance $V_\theta = \tilde{\rho}(1 - \tilde{\rho})$ with subjective posterior belief $\tilde{\rho}$ becomes a function of γ. The function $m(\gamma, V_\theta)$ is then no longer separable in utilities and probabilities. In an extreme case, perfect collinearity between γ and V_θ would render estimates unidentifiable.

3.4 Results

In the following, we discuss the estimates obtained with SIMEX, even though correcting for measurement error has a moderate effect. The SIMEX proce-

[11] In the case of heteroscedastic errors, the variance is inflated by means of contrast vectors (Devanarayan and Stefanski, 2002).

dure yields coefficient estimates that are near the original ones, especially in treatment L, where preferences are measured with smaller errors.

In both treatments, the spread is an increasing concave function of the investment risk V_θ. The estimate β_2 is significantly smaller than 1: $\beta_2 = 0.39$ in NL and $\beta_2 = 0.09$ in L (see Appendix Table C1). Thus, subjects assign disproportionately high weights to assets with low risks, especially in treatment L.

In comparison, risk preferences have a smaller weight ($\beta_1 = 0.81 > \beta_2 = 0.39$ in NL and $\beta_1 = 1.93 > \beta_2 = 0.09$ in L). In both treatments, the investment risk is more important in determining the width of the spread than the inherent risk aversion.

To evaluate Hypothesis 1, we study the correlation between risk preferences and decisions in the two treatments by comparing estimates, marginal effects and predicted values.

Under Hypothesis 1, decisions with and without learning display the same coefficient β_1. Yet, the coefficient β_1 in treatment L is more than twice the coefficient in treatment NL. The difference is illustrated in Figures 3.1a and 3.1b that depict the mean spread in each treatment as a function of the CRRA coefficient γ for low ($V_\theta = 0.05$) and high ($V_\theta = 0.23$) investment risk. After belief updating, mean spreads increase more in risk aversion (dashed lines) compared to the case where no learning takes place (solid lines).

The same effect, albeit attenuated, prevails for the median spread. Figure C4 in Appendix C.3 shows the estimated median spread for risk-averse subjects in treatments NL (solid line) and L (dashed line) in comparison to the EUT benchmark given by decision model (3.4) (dotted line). Note, median spreads in both treatments are higher than in decision model (3.4), but vary less with risk aversion.

The difference in β_1 is not significant, though. The coefficients β_2, on the other hand, significantly differ across the two treatments. The weights (β_1, β_2) in treatment L diverge more from each other than in treatment NL. Because collinearity between preferences and subjective beliefs on the investment risk affects both estimates, it is more appropriate to consider the two factors jointly

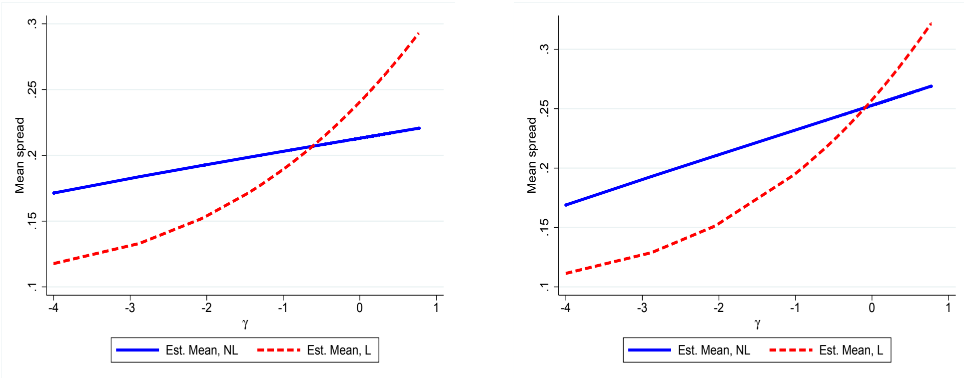

(a) $V_\theta = 0.0475$

(b) $V_\theta = 0.2275$

Figure 3.1: Estimated mean in treatments NL and L for (a) low and (b) high asset variance.

when comparing treatments. I analyze next differences in predicted values and marginal effects.

Table C3 shows the difference in predicted mean spreads with and without learning. Not only differ the treatments with respect to mean spreads, but the difference also increases in risk aversion. An average marginal effect of 0.29 on the difference between predictions casts doubts on Hypothesis 1 (p=0.000 in F-test of zero effect, see Table C3).

As shown in Figures 3.2a and 3.2b, marginal effects on the median spread are significantly lower than marginal effects under the EUT benchmark (dotted line). But more importantly, marginal effects in treatment L are higher than in treatment NL. Without learning, increasing the median risk-averse parameter ($\gamma = 0.32$) by one standard deviation increases the median spread by 1.65% at most (i.e. under the highest variance $V_\theta = 0.2275$ considered in the tables) and has almost no effect on mean spread (maximum increase of 0.10%) (see Appendix Table C2). In contrast, the same effects in treatment L are 3 to 5 times higher for the median, and 40 to 80 times higher for the mean. The difference in marginal effects with and without learning is not significant, though.

In a nutshell, the difference in marginal effects is driven by both the large but non-significant difference in β_1 and the significant difference in β_2. The significant difference in β_2 indicates that subjects' posterior beliefs deviate from

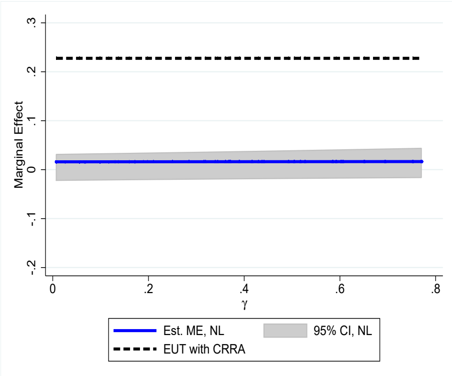

(a) NL

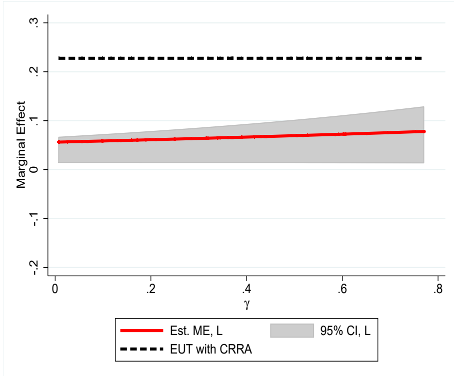

(b) L

Figure 3.2: Marginal effects on the median (with $V_\theta = 0.2275$) in treatments (a) NL and (b) L.

objective probabilities. It is nevertheless not yet clear whether this deviation is unrelated to risk preferences. The importance of risk attitudes after information processing raises the question of whether risk preferences affect investment decisions not only directly through γ, but also through the DM's subjective belief about asset variance $V_\theta(\gamma) = \tilde{\rho}(\gamma)(1 - \tilde{\rho}(\gamma))$. To further explore this question, the following section investigates subjects' inference.

3.5 Testing Bayesian inference

In this section, I analyze to what extent subjects' decisions are consistent with Bayes rule. Furthermore, I examine subjects' probabilistic sophistication separately for different categories of risk attitudes.

Beliefs are not elicited during the investment task. Therefore, I approximate subjects' posterior beliefs using their bid and ask quotes. Under decision model (3.2), the DM's belief corresponds to the midquote, defined as the midpoint of the bid-ask spread: $\frac{a^*+b^*}{2} = \rho$.[12]

Under Bayes rule, the odds for a high asset value given posterior belief ρ sim-

[12]Analogously, midquotes in treatment NL proxy beliefs on marginal probabilities: $\frac{a^*+b^*}{2} = \pi$.

plifies to:

$$\frac{\rho}{1-\rho} = LR(s) \cdot LR(\mu) \tag{3.6}$$

where the likelihood ratio of signal and prior is given by $LR(s) = \frac{q}{1-q}s + \frac{1-q}{q}(1-s)$ and $LR(\mu) = \frac{\mu}{1-\mu}$, respectively.

The DM's sensitivity to the signal and the prior is then estimated with the following model (Grether, 1992; El-Gamal and Grether, 1995):

$$e^{z^*_{ik}} = \left[\frac{\rho}{(1-\rho)}\right]_{ik} = e^{\alpha_0} LR(s_k)^{\alpha_1} LR(\mu_k)^{\alpha_2} e^{u_{ik}} \tag{3.7}$$

with $i = 1, .., N$ subjects and $k = 1, ..., 14$ rounds.

Taking the log gives the equation model for the log odds, z^*_{ik}:

$$z^*_{ik} = \alpha_0 + \alpha_1\, \ell(s_k) + \alpha_2\, \ell(\mu_k) + u_{ik} \tag{3.8}$$

where $\ell(\cdot)$ denotes the log of likelihood ratios.

Bayesian DMs weight the likelihood ratios of the signal and the prior equally as in (3.6).

Therefore, under Bayes rule: $\begin{pmatrix}\alpha_0\\ \alpha_1\\ \alpha_2\end{pmatrix} = \begin{pmatrix}0\\ 1\\ 1\end{pmatrix}$.

The sensitivity of posterior beliefs to risk preferences is analyzed by adding interaction effects into the regression model:

$$\begin{aligned} z^*_{ik} = &\alpha_0 + \alpha_1\, \ell(s_k) + \alpha_1^{RA}\, I_{RA}\, \ell(s_k) + \alpha_1^{RS}\, I_{RS}\, \ell(s_k) + \\ &\alpha_2\, \ell(\mu_k) + \alpha_2^{RA}\, I_{RA}\, \ell(\mu_k) + \alpha_2^{RS}\, I_{RS}\, \ell(\mu_k) + u_{ik} \end{aligned} \tag{3.9}$$

The indicator variables I_{RA} and I_{RS} categorize subjects into risk-averse, and risk-seeking DMs, respectively.[13] The classification is as follows:

$$\begin{cases} (I_{RA}=1)\wedge(I_{RS}=0) & \text{if } \bar{CE}-E(\theta)<-5\%E(\theta) \\ (I_{RA}=0)\wedge(I_{RS}=0) & \text{if } |\bar{CE}-E(\theta)|<5\%E(\theta) \\ (I_{RA}=0)\wedge(I_{RS}=1) & \text{otherwise} \end{cases} \tag{3.10}$$

A subject is classified as risk-neutral ($I_{RA} = I_{RS} = 0$) if her average certainty equivalent deviates from the expected value by less than 5%. Larger negative (positive) deviations define risk-aversion (risk-seeking).

Results in treatment L are compared to the estimates in treatment NL. The log odds in treatment NL take the form: $z^*_{ik} = \alpha_0 + \alpha_3\,\ell(\pi_k) + u_{ik}$, where $\ell(\pi_k)$ denotes the log likelihood ratio of marginal probabilities. Without belief revision, the weight assigned to $LR(\pi)$ significantly deviates from 1 ($\alpha_3 = 0.80$, see Appendix Table C4). That is, subjects conceive the likelihood ratio as biased toward 1. This behavior is consistent with the robust evidence of inverse s-shaped probability weighting functions (see Figure 3.3a).[14]
This inverse s-shape weighting is also found in treatment L, where the likelihood ratio of both the signal and the prior is under-weighted ($\alpha_1 = 0.85$ and $\alpha_2 = 0.88$, see Appendix Table C4). However, this inverse s-shape weighting is less pronounced as both coefficients do not significantly differ from 1.
The effect of risk preferences varies with the information condition. Generally, it stands out that risk attitudes do not interact with information that is provided. They do not affect decision weights in treatment NL. Similarly, they do not interact with given prior beliefs in treatment L.
On average, decisions of risk-neutral subjects are consistent with Bayes rule. The weights assigned to signal and prior information are close to 1 ($\alpha_1 = 1.05$ and $\alpha_2 = 0.95$), as well as the weighting of marginal probabilities ($\alpha_3 = 0.93$). Decisions of risk-seeking subjects also do not deviate from Bayesian inference (($\alpha_1 + \alpha_1^{RS} = 0.95$) and ($\alpha_2 + \alpha_2^{RS} = 0.92$)). Risk-averse subjects, on the other

[13] Classifying subjects into richer categories does not improve the fit as there are no significant differences between moderate and highly risk-averse (risk-seeking) subjects.

[14] Evidence on inverse s-shaped probability weighting can be found in Lattimore et al. (1992); Tversky and Kahneman (1992); Camerer and Ho (1994); Gonzalez and Wu (1999), inter alia.

hand, update their beliefs more conservatively ($\alpha_1 + \alpha_1^{RA} = 0.54$). Their prior belief, too, has a smaller weight ($\alpha_2 + \alpha_2^{RA} = 0.71$) and is thus biased toward the value 0.50, but this deviation is not significant. Figure 3.3b depicts the chosen mid-quotes as a function of objective posterior probabilities. Estimated beliefs of risk-seeking subjects (dotted line) approximate posteriors (dashed line), whereas risk-averse subjects choose, on average, less extreme midquotes. This difference in belief revision becomes more visible in Figures 3.4a and 3.4b that show mean posterior beliefs as a function of priors, separately for risk-seeking and risk-averse subjects. While decisions of risk-seeking subjects (dashed lines) are consistent with Bayesian posteriors, midquotes of risk-averse subjects (solid lines) are closer to priors. Hence, risk-averse subjects update their beliefs more conservatively and, as a result, perceive, on average, a higher investment risk. This finding conforms with the higher correlation between risk-aversion coefficients and chosen spreads in treatment L relative to treatment NL.

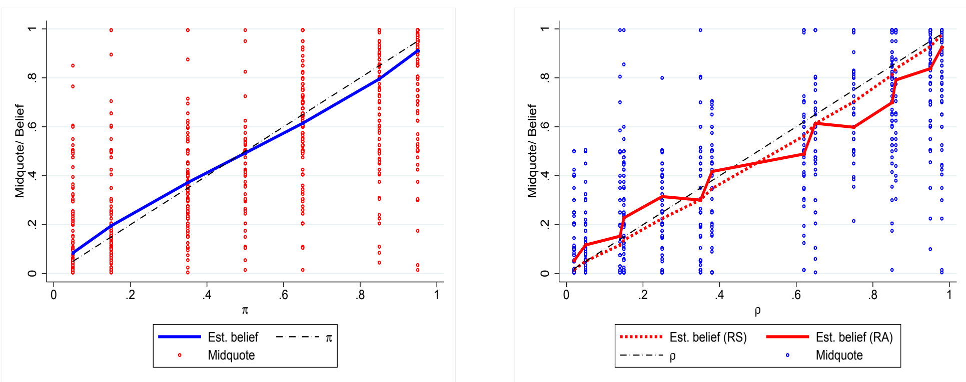

(a) NL **(b)** L

Figure 3.3: Mean subjective belief as a function of objective probabilities in treatments (a) NL and (b) L.

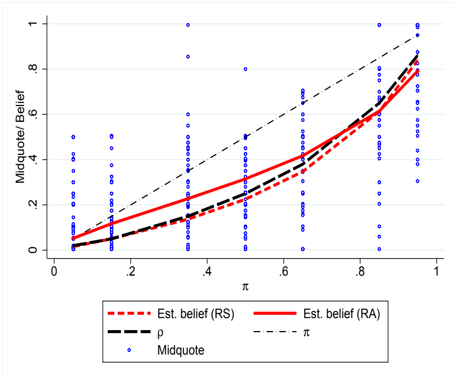

(a) $s = 0$

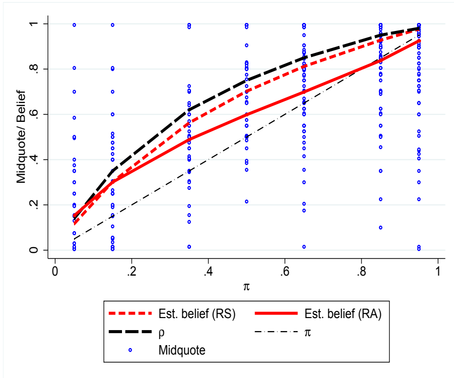

(b) $s = 1$

Figure 3.4: Mean subjective belief as a function of prior probabilities after (a) low and (b) high signals.

3.6 Conclusion

There is extensive experimental literature showing that revealed preferences interact with probabilities. Risky choices are therefore sometimes better modeled with non-expected utility models. However, little is known about how information processing distorts the relation between decisions under risk and objective probabilities.

The analysis in this paper discloses an interaction between information processing and risk attitudes: Learning in this investment setting amplifies the link between decisions and preferences, resulting in a positive correlation between risk aversion and perceived investment risk. A test on Bayes' theorem reveals that risk-averse subjects deviate from Bayesian inference by updating more conservatively. The over-weighting of prior beliefs relative to new evidence is already documented in Edwards (1968); Slovic and Lichtenstein (1971), inter alia. Yet, there is little consensus on the locus of conservatism. Some putative explanations are anchoring on prior information (Slovic and Lichtenstein, 1971), mis-perception or mis-aggregation of information (Edwards, 1968). Alternatively, it might be an aversion toward extreme probabilities (DuCharme, 1970) or a simple heuristic (Gigerenzer and Hoffrage, 1995).

In light of this finding, the question arises whether this observation is driven by a latent effect of cognitive abilities. Unfortunately, the data lacks of cognitive reflection tests. Therefore, the extent to which elicited preferences in this experimental design negatively correlate with cognitive abilities is unknown. Given the salience of expected prize in the preference elicitation task (50 and 75 in the first and second measurement, respectively), it may well be that subjects perceived the task as an assessment of their mathematical skills. In that case, subjects with higher cognitive skills would choose CE closer to the expected prize. However, this argument would not explain the behavior of risk-seeking subjects, whose certainty equivalents deviates from expected prizes but whose inferences conform with Bayes rule.

Despite the necessity of further robustness checks, the results pinpoint to a potential source of conservatism, encouraging future research. A systematic correlation between preferences and belief formation may amplify heterogeneity in decisions, which, in turn, under certain conditions distorts market equilibria.

A Appendix to Chapter 1

A.1 Descriptive statistics

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

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

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

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

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

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

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

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

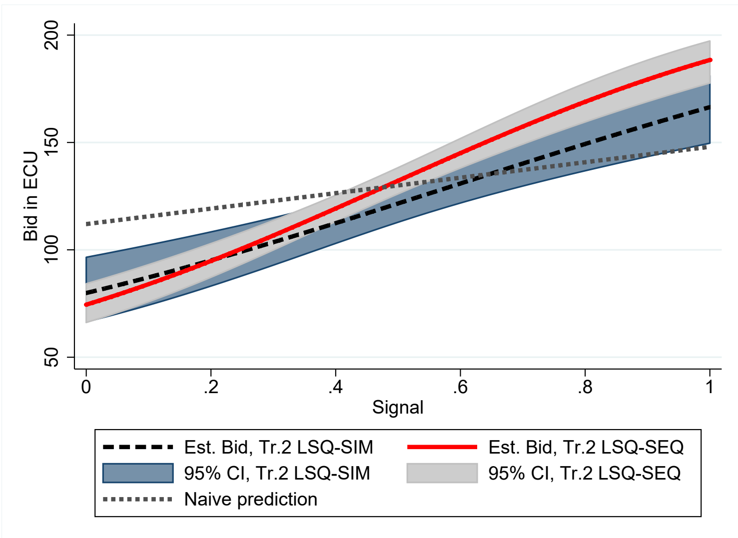

Figure A1: Estimated average bids in treatments LSQ-SIM and LSQ-SEQ.

Table A3: ACTING AGAINST ONE'S OWN SIGNAL (TREATMENT PRICES)

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

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

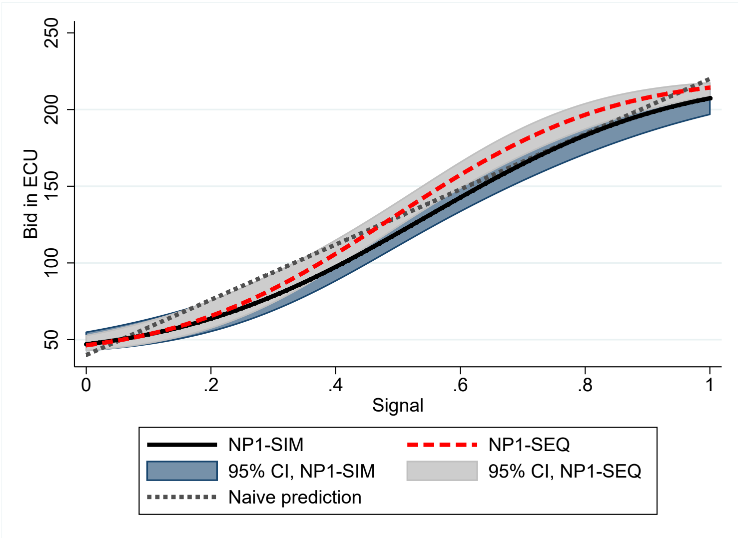

Figure A2: Estimated average bids in treatments NP1-SIM and NP1-SEQ.

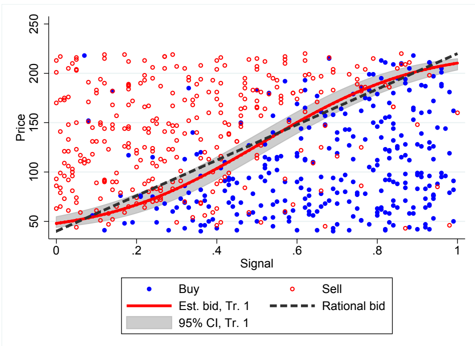

(a) Trader 1

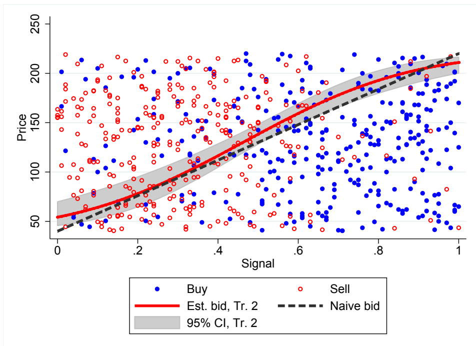

(b) Trader 2

Figure A3: Buys, sells and estimated average bids of traders 1 (a) and 2 (b) in treatment Hyp-SEQ.

Table A4: Shares of naive decisions

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

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

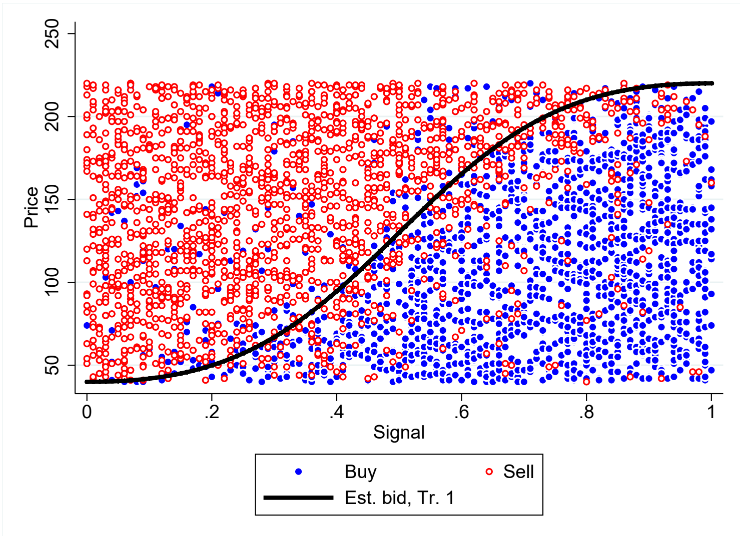

Figure A4: Bid function for trader 1 given random utility model estimates.

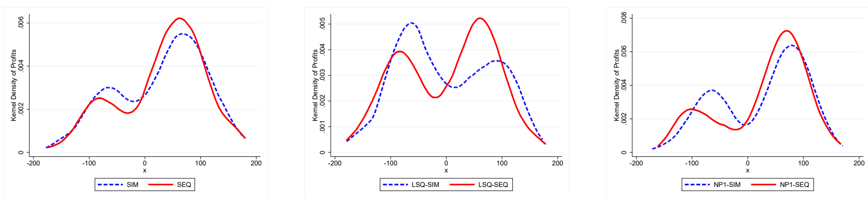

(a) SIM vs. SEQ **(b)** LSQ **(c)** NP1

Figure A5: Kernel density of profits of traders 2 in treatments SIM, SEQ, LSQ-SIM, LSQ-SEQ and NP1-SIM,NP1-SEQ.

Table A5: PROFITS OF TRADERS 2

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

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

A.2 Multiple binomial testing

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

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

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

Table A6: P-VALUES IN MULTIPLE BINOMIAL TESTING

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

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

We account for the multiple testing problem using the Bonferroni significance level of 0.0055 (with a significance level of $\alpha = .05$ for individual tests). Two treatments, SEQ and NP1-SIM, display trading decisions that significantly differ from the naive prediction. In treatment SEQ, the more extreme trading decisions lead to a rejection of the null, while in treatment SIM the share of buys is consistent with naive beliefs. In treatments Hyp-SEQ, NP1-SIM and NP1-SEQ the null is rejected in four out of 9 categories, but only in treatment NP1-SIM the null is rejected after correcting for the multiple testing problem. This significant deviation in the simultaneous mechanism is driven by the overall increased tendency to sell, especially at low signal values. Figure A2 reveals an estimated bidding curve that lies below the naive function for almost all signal values.

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

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

The multiple binomial tests detect in both treatments LSQ-SIM and LSQ-SEQ significant deviations from the share of buys that would be expected under naiveté. The deviations occur at both low and high signal values, reflecting the

Table A7: P-Values in multiple binomial testing (LSQ)

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

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

higher steepness of the bidding curves shown in Figure A1. The information asymmetry helps trader 2 to take into account the price's informativeness.

A.3 Learning

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

Table A8: Proportion of naive decisions

	SIM	SEQ	LSQ-SIM	LSQ-SEQ	Hyp-SEQ	NP1-SIM	NP1-SEQ
First 10	.3971	.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.

A.4 Trading volume

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

Table A9: AVERAGE SIMULATED TRADING VOLUME WITH RANDOM MATCHING OF TRADER 2 PARTICIPANTS

	SIM	SEQ
Main treatments	.8611	.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 $p < .01$.

B Appendix to Chapter 2

B.1 Screen layout

Figures B1a and B1b depict examples of the composition in urn A when the prior is unambiguous and ambiguous, respectively. In the ambiguous urn, the true color of grey balls, which is either red or blue, is unknown.

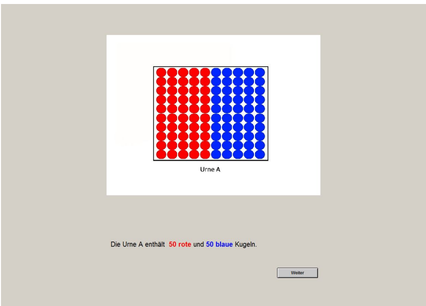

(a) Risky prospect

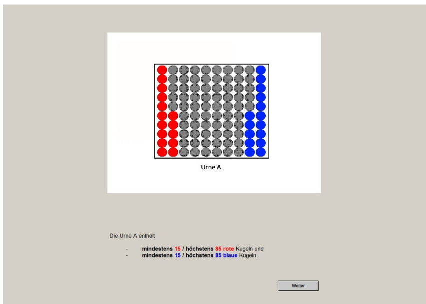

(b) Ambiguous prospect

Figure B1: Examples for visualization of probability distribution with urn A.

In treatment L, a second decision screen is shown to the subjects before they choose their quotes. In the upper left corner, the composition in urn A reminds the subjects of the prior distribution. If the asset takes the value 0 (i.e. the value ball is red), a second ball is drawn from the "urn N". In 75% of all drawings, the subject will then observe a pink ball. The subject will see a green ball with 75% probability if the value ball is blue and the signal ball is drawn from "urn H". The right side of the screen conveys the additional information by showing the color of the signal ball.

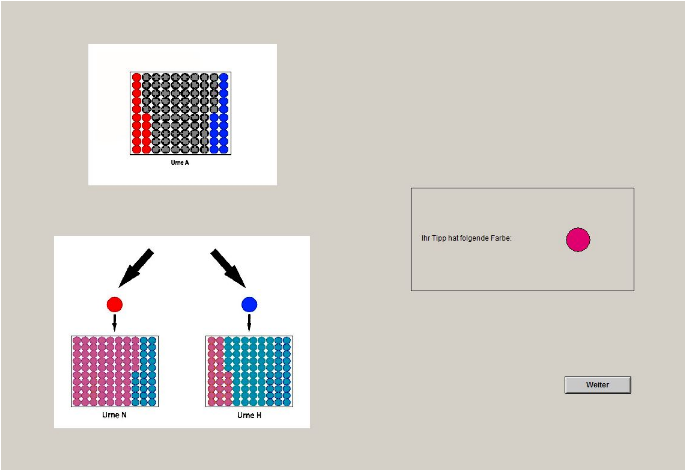

Figure B2: Example for an additional signal at the second stage.

B.2 Bid-ask spread generated by risk aversion

The following derivation shows that risk-aversion introduces a spread between the bid and the ask.

Let $b^{RN} = \mathbb{E}(V)$ be the optimal bid under risk neutrality. Let risk-averse preferences be represented by a strictly concave utility function $U(\cdot)$ with $U'(\cdot) > 0$ and $U''(\cdot) < 0$.
The optimal bid corresponds to the certainty equivalent that makes a risk-averse agent indifferent between the initial position W_0 and the long position. The optimal bid b^{RA} must therefore satisfy:

$$\mathbb{E}_\pi U(W_0 + V - b) = U(W_0)$$

The short-selling ask satisfies accordingly :

$$\mathbb{E}_\pi U(W_0 - V + a) = U(W_0)$$

By Jensen's inequality:

$$\mathbb{E}U(b^{RN}) = \mathbb{E}_\pi U(W_0 + V - b^{RN}) < U(\mathbb{E}_\pi(W_0 + V - E(V)) = U(W_0) = \mathbb{E}U(b^{RA})$$

From $U'(\cdot) > 0$ and $\mathbb{E}U(b^{RN}) < \mathbb{E}U(b^{RA})$, it follows that $b^{RA} < b^{RN} = \mathbb{E}(V)$. Analogously, $a^{RA} > a^{RN} = \mathbb{E}(V)$.

B.3 Results

B.3.1 Reactions to ambiguity

Descriptive statistics

Table B1 shows mean profits for risky and ambiguous prospects, across different ranges of probabilities.

Table B1: Mean profits across different ranges of probabilities

Range of π	[5% – 65%]	[15% – 85%]	[35% – 95%]	Total obs.
Risk	26.82	23.36	23.9	27.97
	(1.82)	(1.88)	(1.85)	(1.52)
Amb.	24.92	18.57	15.63	19.70
	(3.95)	(3.89)	(4.15)	(2.31)
Diff.	1.90	4.79	8.27**	8.27***
	(4.35)	(4.32)	(4.54)	(2.78)
N	840	840	840	1340

Note: *: p-value<.1,**: p-value<.05, ***: p-value<.01. The variable "Amb." represents the indicator variable for the ambiguous rounds.

Table B2 shows median values of bid and ask quotes as a fraction of the expected value. The premia in ambiguous rounds are computed with respect to the midpoint of the probability interval.

Table B2: Median values of quotes as a fraction of the expected value

	$\frac{b}{\mathbb{E}(\pi)}$	$\frac{a}{\mathbb{E}(\pi)}$
Risk	.8	1.0667
Amb.	.6	1.2308
Diff.	-0.20***	-0.1641***

Note: The variable "Amb." represents the dummy variable for the ambiguous rounds. ***: p-value in median test <.01.

Regression estimates

Table B3 presents the results of the median polynomial regression. The estimates for risky prospects are plotted in the Figures 2.3a and 2.3b in Section 2.4.1.

Table B3: Median polynomial regression

Dep. var.	Bid	Ask	Spread
Prior	0.3392***	1.1296***	0.5***
	(.104)	(.104)	(.089)
Prior^2	0.0060***	-0.0018**	-0.005***
	(.001)	(.001)	(.001)
Amb.	-5***	8***	10***
	(1.899)	(2.398)	(3.014)
cons	3.1548**	4.3981**	-1.375
	(1.284)	(2.062)	(.924)
N	1340	1340	1340
R^2	.3717	.3146	.0443

Note: Testing of coeffcients with robust standard errors in parentheses: *: p-value<.1, **: p-value<.05, ***: p-value<.01. The variable "Amb." represents the indicator variable for the ambiguous rounds.

Heterogeneity in ambiguity attitudes

This section examines differences in ambiguity preferences. The behavior in risky rounds is used to identify ambiguity-averse preferences. To control for learning effects, subjects are classified according to their behavior in the last 10 rounds, of which 3 are ambiguous.

In a first step, ambiguous rounds with a set of probabilities $[\pi_l, \pi_h]$ are compared to risky rounds whose probability equals the midpoint $\pi = \frac{\pi_l+\pi_h}{2}$ (henceforth mid-probability). A trading decision is defined as ambiguity-averse (in a tolerant sense) if, for ambiguous prospects, the subject chooses wider spreads than the spread SP_R^{MP} chosen when the unambiguous probability equals the mid-probability. In line with this definition, the variable Y classifies decisions for ambiguous prospects in 3 categories, depending on whether the spread is smaller, equal, or wider than the chosen spread at mid-probabilities:

$$Y_{ij} = \begin{cases} 1 \quad if \quad SP_{\{A,ij\}} > SP^{MP}_{\{R,i\}} \\ 0 \quad if \quad SP_{\{A,ij\}} = SP^{MP}_{\{R,i\}} \\ -1 \quad if \quad SP_{\{A,ij\}} < SP^{MP}_{\{R,i\}} \end{cases}$$

$$i = 1, ..., 67 \texttt{ subjects}, j = 1, 2, 3 \texttt{ ambiguous decisions}$$

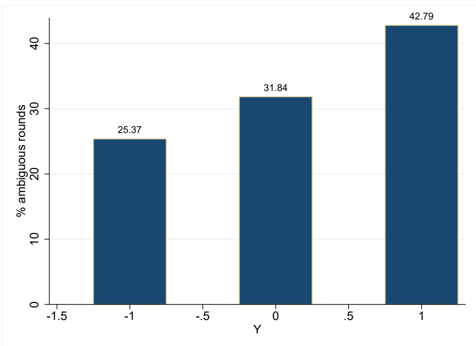

(a) Y

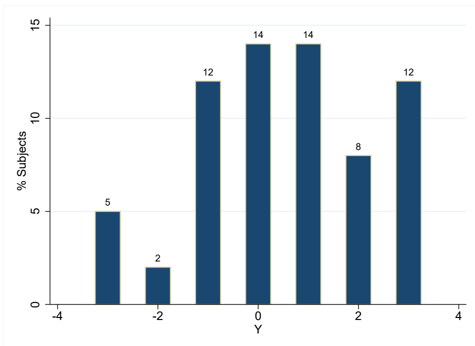

(b) Sum of Y

Figure B3: Classification of ambiguous decisions (a) and subjects (b) in the last 10 rounds.

Subjects exhibit a larger spread in 43% of the ambiguous rounds (see Figure

B3a). Out of the 67 subjects, 12 chose a wider spread for all 3 ambiguous rounds versus 5 subjects who always chose a smaller spread (see Figure B3b). The majority of subjects is, on average, ambiguity-averse with $\sum_{j=1}^{3} Y_i > 0$.

A second step differentiates more thoroughly between ambiguity aversion and subjective expected utility (SEU). For instance, when $\pi \in [.05, .65]$, a risk-averse agent with subjective belief $\Pi^* = .45$ chooses a wider spread than for a risky prospect with $\pi = .35$, without being ambiguity-averse. For this purpose, decisions with $Y = 1$ are further distinguished into decisions that can and cannot be explained by SEU as well. A trading strategy is defined to be inconsistent with SEU when the spread for an ambiguous prospect with $\pi \in [\pi_l, \pi_h]$ is wider than *any* chosen spread for all risky prospects with $\pi \in [\pi_l, \pi_h]$. For instance, the subject's decision in an ambiguous round with $\pi \in [.05, .65]$ is compared to all his decisions in risky rounds with $\pi \in \{.05, .15, .35, .50, .65\}$.[1] In treatment NL, 52.33% of the rounds that are consistent with ambiguity-averse preferences reject SEU as a possible explanation.

In treatment L, significantly less ambiguous rounds can be classified as consistent with ambiguity-averse preferences (33% in L vs. 43% in NL, compare Figures B3a & B4). This is mainly due to some large spreads in risky rounds. In these rounds, contradicting signals endogenously create uncertainty: subjects choose wider spreads for risky prospects when the signal contradicts prior beliefs. Still, out of these ambiguous decisions with wider spreads, 52.31% display a spread wider than any chosen spread for unambiguous priors within $[.15, .85]$, and exclude therefore SEU.

B.3.2 Learning

Figure B5a depicts the distribution of chosen spreads in the ambiguous rounds of treatment NL with $\pi \in [.05, .65]$ or $[.35; .95]$. Figure B5b refers to the dis-

[1] Theoretically, subjects should choose a maximum spread at a prior at 50%, where the lottery exhibits the highest risk. Nevertheless, in treatment NL 27 subjects choose a maximum spread at a different probability. Therefore, spreads are compared to the maximum spread that each subject has chosen for the risky prospects, regardless of whether it has been chosen at a probability of 50% or not.

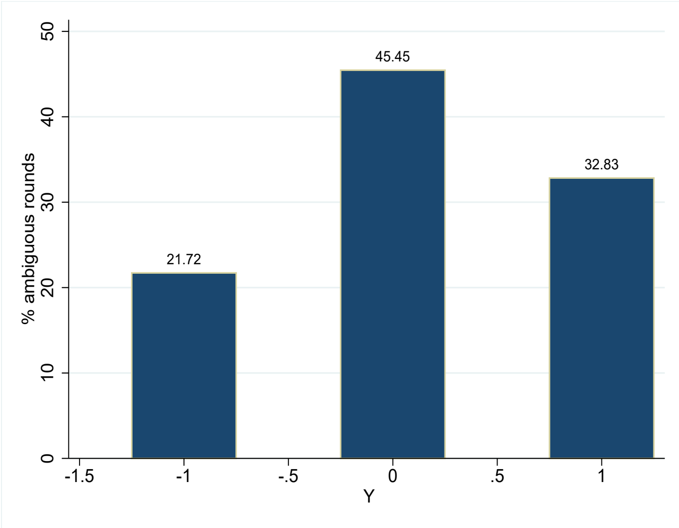

Figure B4: Percentage of ambiguous rounds with smaller, equal and larger spreads in AI .

tribution of spreads in the ambiguous rounds of treatment L with $\rho \in [.05, .65]$ or $[.35; .95]$. In both figures, the vertical solid and dashed lines represent the median and mean spread, respectively. The distributions of spreads do not differ for the same range of marginal and posterior probabilities (p-value=.92 in Kolmogorov-Smirnov test).

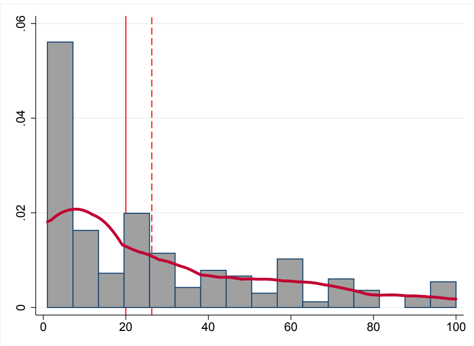

(a) Ambiguous rounds in NL

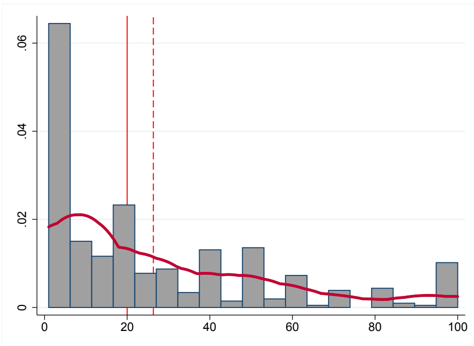

(b) Ambiguous rounds in L

Figure B5: Spreads for ambiguous prospects for the same theoretical dispersion in marginal (a) and Bayesian posterior (b) probabilities.

Results of NNLS-SUR

Table B4 shows the coefficient estimates of the NNLS-SUR model. Because, by design, there is less variation in the ambiguous probabilities, the estimation is more robust when assuming symmetric premia in the bid and the ask. The model estimates assuming Bayesian update of recursive preferences (BRP) do not differ from the estimates in treatment NL (p-value of 1 for the ask equation and .2211 for the bid equation in the Lagrange-Multiplier test).

Table B4: Coefficient estimates for probability weighting function and risk premia

	NL	L		
	mid-prior	FBU***	MLU***	BRP
β	0.9646	1.1532	0.9206	0.9824
	(.0370)	(.0301)	(.0440)	(.0493)
α	0.6563	1.1754	0.2574	0.6658
	(.0754)	(.0686)	(.0288)	(.0744)
RP	0.2982	0.2491	0.2491	0.2491
	(.0258)	(.0301)	(.0301)	(.0301)

Note: Nonlinear least squares estimation with CRSE in a seemingly unrelated regression. ***: p-value<.01, refers to a significance difference between the model estimates in treatment NL and the ones with updated beliefs in Lagrange-Multiplier tests.

Heterogeneity in updating

Figure B6 depicts mid-quotes for risky prospects and the mean regression estimates as a function of unambiguous priors. The dashed and solid lines correspond to mean estimates after subjects receive a high and a low signal, respectively. The average mid-quote increases in the prior, showing no evidence of base-rate-neglect.

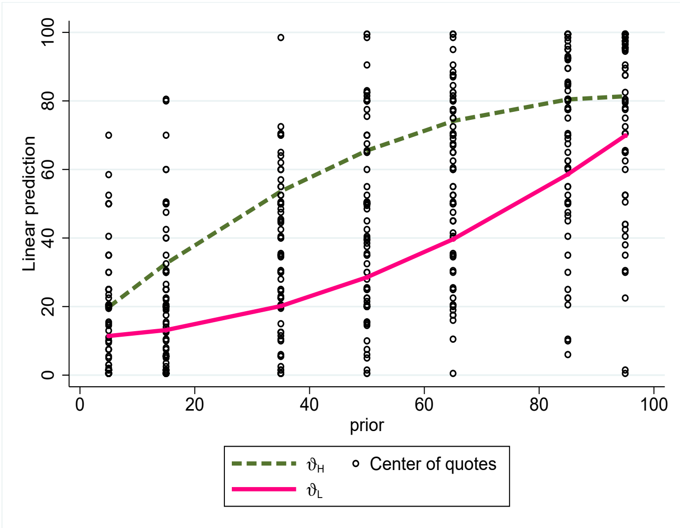

Figure B6: Mid-quotes for unambiguous assets and their mean-estimates for the two signals and the group of Bayesian updaters (clusters 3 and 5).

B.3.3 Cluster analysis

To discern the different ranges of updated beliefs and their prevalence, bid-ask pairs for ambiguous prospects are clustered. The cluster analysis is performed in k-medians with 8 clusters, yielding the 8 different ranges for updated beliefs listed in Table B5.[2]

In total, 21.72% of the ambiguous decisions belong to the clusters 1 and 6 and are consistent with MLU. Quotes in these clusters are close to one extremum and exhibit, on average, the smallest spread of 1 ECU. The opposite behavior is described in clusters 7 and 8, that represent 25.75% of the bid-ask pairs. These observations exhibit a substantial spread of more than 30 ECU. In approximately one third of these decisions, the spread is chosen wide enough to implement almost surely a no-trade outcome (cluster 8). In cluster 4, 10.35% of

[2] The value of 8 clusters finds its justification in the theory, allowing the identification of 8 clusters in the upper triangular grid of bid-ask pairs: extreme beliefs upon both a low and a high signal (centered around the bid-ask points: (0,0); (100,100)), ambiguity-neutral Bayesian beliefs upon both a low and a high signal (the 45 line (5,5) to (95,95)), ambiguity-averse Bayesian beliefs upon both a low and a high signal ((5,65), (35,95)), maximum ambiguity-aversion (0,100), ambiguity-neutral likelihood-insensitive beliefs (50,50). Robustness checks with more and less clusters do not yield a better comprehension of the data.

Table B5: Median bids, asks and spreads and corresponding statistics for 8 clusters in ambiguous rounds of treatment L

Cluster	bid	ask	spread	% trade	% obs	consistent with
1	1	2	1	100	12.12	MLU
2	10	16	5	100	5.30	AN & Bayesian
3	15	33.5	20	86.36	16.67	Bayesian
4	40	50	5	95.12	10.35	AN-LI/conservatism
5	60	80	20	82.5	20.20	Bayesian
6	98.5	99	1	100	9.60	MLU
7	20	70	50	53.62	17.42	AA
8	1	99	98	15.15	8.33	AA - non-participants

Note: Cluster analysis in k-medians.

the quotes disclose a small spread with bids and asks around 50%, the midpoint of the set of priors. These quotes match the behavior of an ambiguity-neutral but likelihood-insensitive (AN-LI) investor who is rather unresponsive to incoming information. Under the assumption that subjects have second-order beliefs, whose mean equals the midpoint of the set of priors, over-emphasizing the mid-prior 50% concurs with conservatism. Conservatism predicts an over-weighting of the prior belief, but no increase in the spread. The remainder of the decisions amounts to 42.17% of bid-ask pairs in the clusters 2, 3 and 5. These quotes are consistent with Bayesian updating. The decisions in cluster 2 result in small spreads and, compared to decisions under risk, do not show any evidence of ambiguity aversion. The majority of the bid-ask pairs, though, falls in clusters 3 or 5 that disclose a median spread of 20 ECU. Figure B7 summarizes the results of the cluster analysis. It depicts bid-ask pairs that are consistent with MLU, ambiguity-averse Bayesian, ambiguity-averse non-Bayesian and ambiguity-neutral beliefs in diamonds, squares, triangles and dots or crosses, respectively.

Table B6 lists the results of the same cluster analysis in treatment NL. The analysis yields less extreme clusters of beliefs. Furthermore, the observations are

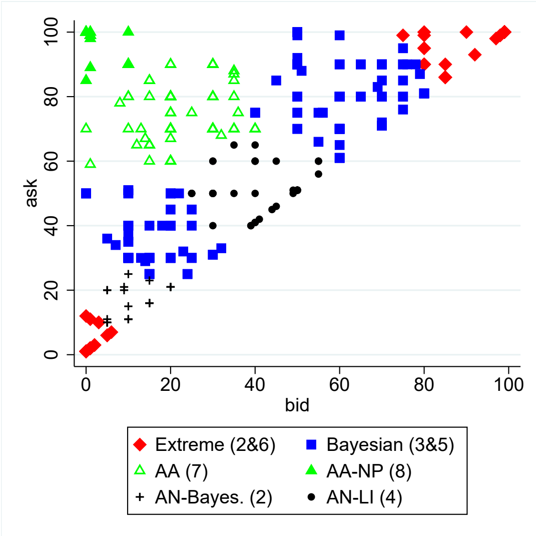

Note to the legend: Cluster categories in parentheses.

Figure B7: Clusters of bid-ask pairs in ambiguous rounds of treatment L

distributed more evenly across the eight clusters, yielding the more symmetric distributions of quotes reflected in Figures 2.6a and 2.6b.

Table B6: Median bids, asks and spreads and corresponding statistics for 8 clusters in ambiguous rounds of treatment NL

Cluster	bid	ask	spread	% trade	% obs
1	5	10.5	1	100	8.96
2	20	25	2	95	7.46
3	30	37	1	97.06	12.69
4	20	50	35	69.70	12.31
5	49	60	12.5	82.61	17.46
6	70	80	5	90.70	16.04
7	35	90	52.5	34.38	11.94
8	4.5	85	72.5	19.44	13.43

Note: Cluster analysis in k-medians.

C Appendix to Chapter 3

C.1 Descriptive statistics and design

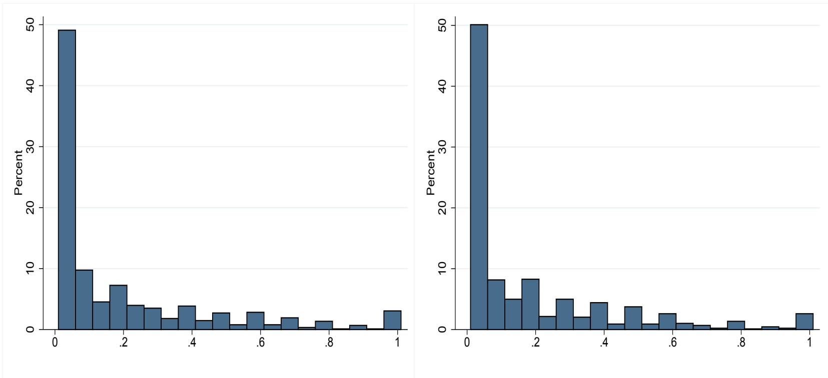

(a) Spread in NL **(b)** Spread in L

Figure C1: Distribution of chosen spreads in treatments (a) NL and (b) L.

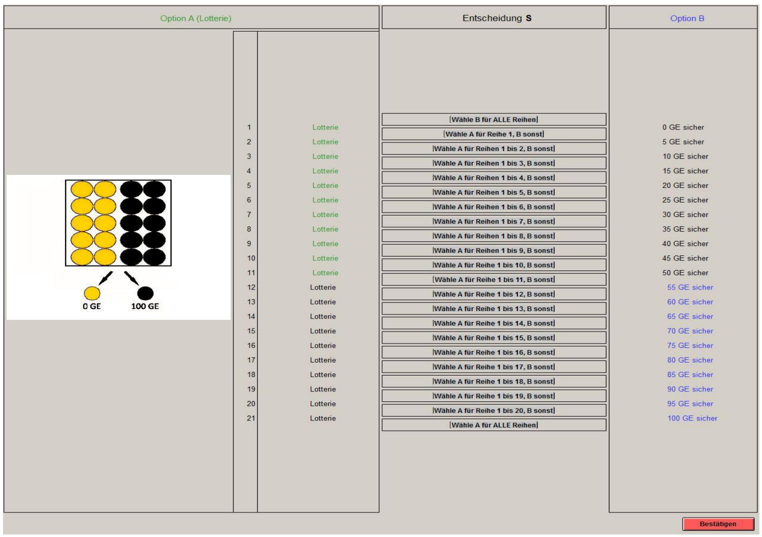

Figure C2: Example of computer interface in Part 2

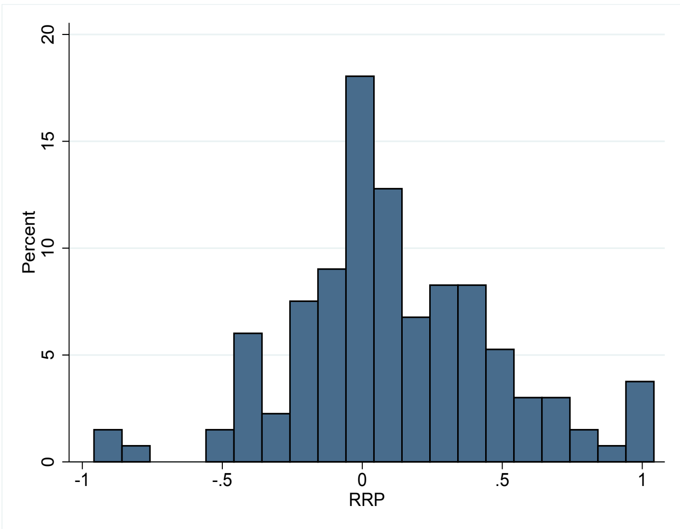

Figure C3: Distribution of relative risk premia

C.2 Approximation to optimal bid and ask under CRRA

Define the lottery in the asset's value as $\tilde{\theta} \equiv (1, \pi; 0)$. The optimal bid b and (short-selling) ask a satisfy respectively:

$$EU(W_0 + \tilde{\theta} - b) = U(W_0) \tag{C.1}$$

$$EU(W_0 - \tilde{\theta} + a) = U(W_0) \tag{C.2}$$

Re-writing the lottery with a zero-mean risk yields: $\tilde{\theta} = E[\tilde{\theta}] + \tilde{x} = \mu + \tilde{x}$, where $\mu = \pi$ and $\tilde{x} \equiv ((1-\pi), \pi; -\pi, (1-\pi))$ with $E\tilde{x} = 0$.

Let k be a risk inflating factor. Define $\Gamma(k)$ as the risk premium in the bid and the ask. Following Gollier (2001), we can rewrite Equations C.1 and C.2 as:

$$EU(W_0 + k(\mu + \tilde{x}) - (k\mu - \Gamma(k))) = U(W_0) \tag{C.3}$$

$$EU(W_0 - k(\mu + \tilde{x}) + (k\mu + \Gamma(k))) = U(W_0)$$

For illustration, we derive the Arrow-Pratt approximation for the bid.

The first derivative of Equation C.3 with respect to k yields:

$$EU'(W_0 + \tilde{x} + \Gamma(k))[\tilde{x} + \Gamma'(k)] = 0 \tag{C.4}$$

Thus,

$$\Gamma'(k) = -\frac{E\tilde{x}U'(\cdot}{EU'(\cdot)} \tag{C.5}$$

Assuming that without risk, i.e. with $k = 0$, $\Gamma(k) = 0$ and because $E\tilde{x} = 0$:

$$\Gamma'(0) = 0$$

The second derivative of Equation C.3 with respect to k yields:

$$EU''(W_0 + \tilde{x} + \Gamma(k))[\tilde{x} + \Gamma'(k)]^2 + \Gamma''(k)EU'(W_0 + \tilde{x} + \Gamma(k)) = 0 \tag{C.6}$$

Thus,

$$\Gamma''(k) = -\frac{EU''(\cdot)[\tilde{x} + \Gamma'(k)]^2}{EU'(\cdot)} \tag{C.7}$$

and

$$\Gamma''(0) = -\frac{U''(W_0)}{U'(W_0)}E\tilde{x}^2 = A(W_0)Var[\tilde{x}] \tag{C.8}$$

The term $A(W_0)$ corresponds to the absolute coefficient of risk aversion.

The Taylor approximation around $k = 0$ corresponds to:

$$\Gamma(k) \approx \Gamma(0) + k\Gamma'(0) + \frac{1}{2}k^2\Gamma''(0) \tag{C.9}$$

For $k = 1$, we obtain:

$$\Gamma(1) \approx \frac{1}{2}A(W_0)Var[\tilde{x}] \tag{C.10}$$

The Arrow-Pratt approximations to the optimal bid and ask are accordingly:

$$b = \mu - \frac{1}{2}A(W_0)Var[\tilde{x}] \tag{C.11}$$

$$a = \mu + \frac{1}{2}A(W_0)Var[\tilde{x}] \tag{C.12}$$

For a utility function from the class of CRRA, $A(W_0) = \gamma W_0^{-1}$:

$$b = \pi - \frac{1}{2}\gamma W_0^{-1}Var[\tilde{x}] \tag{C.13}$$

$$a = \pi + \frac{1}{2}\gamma W_0^{-1}Var[\tilde{x}] \tag{C.14}$$

where $Var[\tilde{x}] = \pi(1-\pi)$. Inserting $W_0 = 1$ yields a^* and b^*:

$$b^* = \pi - \frac{1}{2}\gamma Var[\tilde{x}] \tag{C.15}$$

$$a^* = \pi + \frac{1}{2}\gamma Var[\tilde{x}] \tag{C.16}$$

The spread corresponds then to:

$$a^* - b^* = \gamma Var[\tilde{x}] \tag{C.17}$$

C.3 Estimates of the Tobit model

Table C1: Maximum likelihood estimates of two limit non-linear Tobit model under risk

	No Learning		Learning		Diff:(L-NL)
	Risk	Risk^{SIMEX}	Risk	Risk^{SIMEX}	Risk^{SIMEX}
β_0	-2.0698***	-2.0762***	-2.0521***	-2.0495***	0.0567
	(0.152)	(0.168)	(0.116)	(0.118)	(0.205)
β_1	0.6491	0.8145	1.9706	1.9301	1.1156
	(0.621)	(0.785)	(0.730)	(0.829)	(1.141)
β_2	0.4185***	0.3917***	0.0890***	0.0936***	-0.2981***
	(0.219)	(0.260)	(0.082)	(0.088)	(0.274)
σ	1.0517***	1.0517***	0.8864***	0.8888***	-0.1629**
	(0.041)	(0.041)	(0.047)	(0.047)	(0.062)
LL	-169.41		-207.76		
N	938	938×10^3	882	882×10^3	

Note: $^{*}: p < 0.1,^{**}: p < 0.05,^{***}: p < 0.01$ The SIMEX-coefficients were obtained with 1000 re-measurements of predictors. Standard errors in parentheses are clustered at subject level. Tests for coefficients $(\beta_0, \beta_1, \beta_1)$ refer to deviation from $(\beta_0, \beta_1, \beta_2)' = (0, 1, 1)'$.

The SIMEX procedure with heteroscedastic errors and replicate measures consists of the following steps (cf. Devanarayan and Stefanski 2002; Carroll et al. 2006):

1. A normalized contrast vector is created:

$$c_{b,i,j} = \frac{z_{b,i,j} - \bar{z}_{b,i,.}}{\sqrt{\sum(z_{b,i,j} - \bar{z}_{b,i,.})^2}}$$

   with $z_{b,i,j} \sim N(0,1)$.

2. Error-contaminated risk measures are generated with inflation factor $\zeta \in [0, 1.6]$:

$$\tilde{\gamma}_{b,i}(\zeta) = \gamma_i + (\zeta/2)^{1/2} \sum_{j=1,2} c_{b,i,j}\gamma_{i,j}$$

3. The Tobit model is estimated yielding $\beta(\zeta)$.

4. The procedure is repeated for B=1000 variance-inflated data sets.

5. A link function is estimated using a quadratic extrapolant function: $E[\beta_b|\zeta] = \delta_0 + \delta_1\zeta + \delta_2\zeta^2, \quad b = 1, ..., 1000.$

6. SIMEX estimates are given by extrapolation to $\zeta = -1$, i.e. $\beta^{SIMEX} = E[\beta_b|\zeta = -1]$.

Because the regressor γ is transformed, the size of the coefficients is not informative. I interpret instead changes in the mean and median. The Equations C.18 and C.19 show the computation of the expected mean (cf. Amemiya 1985; Greene 2008; Carson and Sun 2007).
For simplification, define $m(0, V_\theta) =: \underline{m}$ and $m(\gamma, V_L) =: m$. The function $\Phi(\cdot)$ denotes the standard normal cumulative distribution function.

$$\begin{aligned} E[y^*|(\underline{y}^* < y^* < 1), \gamma, V_\theta] &= E[\exp\left(m(\gamma, V_L) + \epsilon\right)|\underline{m} < m + \epsilon < 0] \\ &= \exp(m)E[\exp(\epsilon)|\underline{m} - m < \epsilon < -m] \\ &= \left[\Phi(\frac{-m}{\sigma}) - \Phi(\frac{\underline{m} - m}{\sigma})\right]^{-1} \exp(m) \\ &\int_{\underline{m}-m}^{-m} \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left(- \frac{(\epsilon - \sigma^2)^2}{2\sigma^2} + \frac{\sigma^4}{2\sigma^2}\right) d\epsilon \\ &= \left[\Phi(\frac{-m}{\sigma}) - \Phi(\frac{\underline{m} - m}{\sigma})\right]^{-1} \exp\left(m + \frac{\sigma^2}{2}\right) \\ &\left[\Phi(\frac{-m - \sigma^2}{\sigma}) - \Phi(\frac{\underline{m} - m - \sigma^2}{\sigma})\right] \end{aligned} \tag{C.18}$$

$$\begin{aligned} E[y|\gamma, V_\theta] &= 0 \cdot Pr[y^* \leq \underline{y}^*|\gamma, V_\theta] + 1 \cdot Pr[y^* \geq 1|\gamma, V_\theta] \\ &+ E[y^*|(\underline{y}^* < y^* < 1), \gamma, V_\theta] Pr[\underline{y}^* < y^* < 1|\gamma, V_\theta] \\ &= 1 \cdot \Phi\left(- \frac{m}{\sigma}\right) + \exp\left(m + \frac{\sigma^2}{2}\right) \\ &\left[\Phi(\frac{-m - \sigma^2}{\sigma}) - \Phi(\frac{\underline{m} - m - \sigma^2}{\sigma})\right] \end{aligned} \tag{C.19}$$

Accordingly, given $(\underline{y}^* < y^* < 1)$ and assuming the conditional median of ϵ given (γ, V_θ) to be zero, the median spread equals $exp(m)\mathbb{1}_{\{\underline{m}<m<0\}}$.

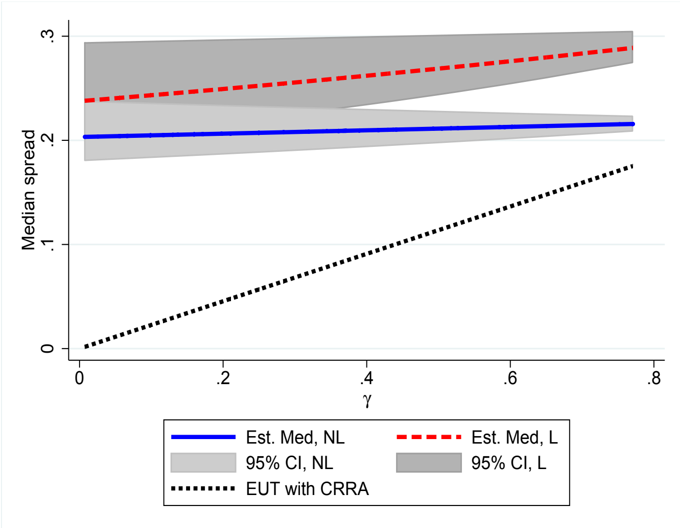

Figure C4: Estimated median spread for $\gamma > 0$ and $V_\theta = 0.2275$

Table C2: Moment estimates and change in moments

	$V_\theta = .0475$	$V_\theta = .1275$	$V_\theta = .2275$
Risk - No Learning			
$E[y\|(\underline{y}^* < y^* < 1), \bar{\gamma}, V_\theta]$	0.7914	0.7953	0.7986
$\%\Delta(E[y\|(\underline{y}^* < y^* < 1), \bar{\gamma}, V_\theta], \bar{\gamma} + \bar{\hat{\sigma}})$	0.04	0.07	0.10
Median$(y\|\bar{\gamma}, V_\theta)$	0.1648	0.1875	0.2077
$\%\Delta(median(y\|\bar{\gamma}, V_\theta), \bar{\gamma} + \bar{\hat{\sigma}})$	0.89	1.31	1.65
Risk- Learning			
$E[y\|(\underline{y}^* < y^* < 1), \bar{\gamma}, V_\theta]$	0.8334	0.8360	0.8376
$\%\Delta(E[y\|(\underline{y}^* < y^* < 1), \bar{\gamma}, V_\theta], \bar{\gamma} + \bar{\hat{\sigma}})$	3.16	3.64	3.97
Median$(y\|\bar{\gamma}, V_\theta)$	0.2436	0.2591	0.2694
$\%\Delta(median(y\|\bar{\gamma}, V_\theta), \bar{\gamma} + \bar{\hat{\sigma}})$	4.66	5.11	5.41

Note: Consistent with values in Table 3.1: $\bar{\gamma} = 0.32, \bar{\hat{\sigma}} = 0.21$.

Table C3: Regression model of differences in predictions

	Diff.
γ	-0.2323***
	(0.004)
γ^2	0.3268***
	(0.004)
V_θ	-0.1250***
	(0.001)
N	1820
R^2	0.9943
AME_γ	0.2972
MEM_γ	0.3212

Note: $^{*}: p < 0.1, ^{**}: p < 0.05, ^{***}: p < 0.01$. Robust standard errors clustered at subject level (CRSE) in parentheses. AME and MEM denote average marginal effect and marginal effect at the median coefficient $\gamma = 0.05$, respectively.)

C.4 Bayesian inference

Table C4: Regression estimates for Bayesian inference

	Learning		No Learning	
cons	-0.1871**	-0.1896**	-0.0285	-0.0285
	(0.077)	(0.078)	(0.077)	(0.078)
$\ell(s)$	0.8585	1.0463		
	(0.118)	(0.226)		
$\ell(\mu)$	0.8836	0.9501		
	(0.080)	(0.132)		
$\ell(\pi)$			0.7995***	0.9262***
			(0.058)	(0.076)
$I_{RA} \times \ell(s)$		-0.5111**		
		(0.257)		
$I_{RS} \times \ell(s)$		-0.0936		
		(0.308)		
$I_{RA} \times \ell(\mu)$		-0.2336		
		(0.190)		
$I_{RS} \times \ell(\mu)$		-0.0274		
		(0.188)		
$I_{RA} \times \ell(\pi)$				-0.1468
				(0.114)
$I_{RS} \times \ell(\pi)$				-0.2822
				(0.172)
R^2	0.53	0.55	0.51	0.52
N	882	882	938	938

Note: $^{*}: p < 0.1, ^{**}: p < 0.05, ^{***}: p < 0.01$. Robust standard errors clustered at subject level (CRSE) in parentheses.

Bibliography

Abdellaoui, M., Baillon, A., Placido, L., and Wakker, P. P. The Rich Domain of Uncertainty: Source Functions and Their Experimental Implementation. *The American Economic Review*, 101(2):695–723, 2011.

Ahn, D., Choi, S., Gale, D., and Kariv, S. Estimating ambiguity aversion in a portfolio choice experiment. *Quantitative Economics*, 5:195–223, 2014.

Allais, M. L'Extension des Theories de l'Equilibre Economique General et du Rendement Social au Cas du Risque. *Econometrica*, 21(2):269–290, 1953.

Amemiya, T. *Advanced econometrics*. Harvard University Press, 1985. ISBN 9780674005600.

Amihud, Y., Mendelson, R., and Lauterbach, B. Market microstructure and securities values: Evidence from the Tel Aviv Stock Exchange. *Journal of Financial Economics*, 45:365–390, 1997.

Andersen, S., Harrison, G. W., Lau, M. I., and Rutström, E. E. Elicitation using multiple price list formats. *Experimental Economics*, 9(4):383–405, dec 2006.

Antoniou, C., Harris, R. D. F., and Zhang, R. Ambiguity Aversion and Stock Market Participation: An Empirical Analysis. 2015.

Baillon, A., Bleichrodt, H., Keskin, U., Haridon, O. L., and Li, C. Learning under ambiguity: An experiment using initial public offerings on a stock market. 2013.

Baillon, A., Halevy, Y., and Li, C. Experimental Elicitation of Ambiguity Attitude using the Random Incentive System. 2015.

Ballinger, P. T. and Wilcox, N. T. Decisions, Error and Heterogeneity. *The Economic Journal*, 107(443):1090–1105, 1997.

Battalio, R. C., Kagel, J. H., and Jiranyakul, K. Testing Between Alternative Models of Choice Under Uncertainty: Some Initial Results. *Journal of Risk and Uncertainty*, 3(1):25–50, 1990.

Bazerman, M. H. and Samuelson, W. F. I Won the Auction But Don't Want the Prize. *Journal of Conflict Resolution*, 27(4):618–634, 1983.

Bell, D. E. Regret in Decision Making under Uncertainty. *Operations Research*, 33(1):1–27, 1985.

Bernasconi, M. and Loomes, G. Failures of the reduction principle in an Ellsberg-type problem. *Theory and Decision*, 32(1):77–100, 1992.

Bossaerts, P., Ghirardato, P., Guarnaschelli, S., and Zame, W. R. Ambiguity in Asset Markets: Theory and Experiment. *The Review of Financial Studies*, 23(4):1325–1359, 2010.

Brown, D. P. and Zhang, Z. M. Market Orders and Market Efficiency. *The Journal of Finance*, 52(1):277–308, 1997.

Bruhin, A., Fehr-Duda, H., and Epper, T. Risk and Rationality: Uncovering Heterogeneity in Probability Distribution. *Econometrica*, 78(4):1375–1412, 2010.

Camerer, C. and Weber, M. Recent developments in modeling preferences under risk. *Journal of Risk and Uncertainty*, (5):325–370, 1992.

Camerer, C. F. An experimental test of several generalized utility theories. *Journal of Risk and Uncertainty*, 2(1):61–104, 1989.

Camerer, C. F. and Ho, T. H. Violations of the betweenness axiom and nonlinearity in probability. *Journal of Risk and Uncertainty*, 8(2):167–196, 1994.

Cao, H. H., Wang, T., and Zhang, H. H. Model Uncertainty, Limited Market Participation , and Asset Prices. *The Review of Financial Studies*, 18(4): 1219–1251, 2005.

Carrillo, J. D. and Palfrey, T. R. No trade. *Games and Economic Behavior*, 71 (1):66–87, jan 2011.

Carroll, R. J., Ruppert, D., Stefanski, L. A., and Crainiceanu, C. M. *Measurement Error in Nonlinear Models: A Modern Perspective*. 2006. ISBN 9781584886334.

Carson, R. T. and Sun, Y. The Tobit model with a non-zero threshold. *Econometrics Journal*, 10(3):488–502, 2007.

Chapman, D. A. and Polkovnichenko, V. First-Order Risk Aversion, Heterogeneity, and Asset Market Outcomes. *The Journal of Finance*, 64(4):1863–1887, 2009.

Charness, B. G. and Levin, D. The Origin of the Winner's Curse: A Laboratory Study. *American Economic Journal: Microeconomics*, 1(1):207–236, 2009.

Chateauneuf, A., Eichberger, J., and Grant, S. Choice under uncertainty with the best and worst in mind: Neo-additive capacities. *Journal of Economic Theory*, 137(1):538–567, 2007.

Chow, C. and Sarin, R. K. Known, unknown, and unknowable uncertainties. *Theory and Decision*, 52(2):127–138, 2002.

Cohen, M., Gilboa, I., Ja, J. Y., and Schmeidler, D. An experimental study of updating ambiguous beliefs. *Risk, Decision and Policy*, 5(2):123–133, 2000.

Comerton-Forde, C., Malinova, K., and Park, A. Regulating Dark Trading: Order Flow Segmentation and Market Quality. 2016.

Conlisk, B. J. Three Variants on the Allais Example. *The American Economic Review*, 79(3):392–407, 1989.

Conte, A., Hey, J. D., and Moffatt, P. G. Mixture models of choice under risk. *Journal of Econometrics*, 162(1):19–88, 2011.

Cook, J. R. and Stefanski, L. A. Simulation-Extrapolation Estimation in Parametric Measurement Error Models. *Journal of the American Statistical Association*, 89(428):1314–1328, 1994.

Crawford, V. P. and Iriberri, N. 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.

De Filippis, R., Guarino, A., Jehiel, P., and Kitagawa, T. Updating Ambiguous Beliefs in a Social Learning Experiment. 2016.

De Jong, F. and Rindi, B. *The Microstructure of Financial Markets*. Cambridge edition, 2009.

Devanarayan, V. and Stefanski, L. A. Empirical simulation extrapolation for measurement error models with replicate measurements. *Statistics and Probability Letters*, 59:219–225, 2002.

Diamond, D. W. and Verrecchia, R. E. Information Aggregation in a Noisy Rational Expectations Economy. *Journal of Financial Economics*, 9:221–235, 1981.

Dimmock, S. G., Kouwenberg, R., and Wakker, P. P. Ambiguity Attitudes in a Large Representative Sample. *Management Science*, 62(5):1363–1380, 2016.

Dow, J. and Werlang, S. R. d. C. Uncertainty Aversion, Risk Aversion, and the Optimal Choice of Portfolio. *Econometrica*, 60(1):197–204, 1992a.

Dow, J. and Werlang, S. R. d. C. Excess volatility of stock prices and Knightian uncertainty. *European Economic Review*, 36(2-3):631–638, 1992b.

DuCharme, W. M. Response Bias Explanation of Conservative Human Inference. *Journal of Experimental Psychology*, 85(1):66–74, 1970.

Easley, D. and Hara, M. O. Microstructure and Ambiguity. *The Journal of Finance*, 65(5):1817–1846, 2010.

Edwards, W. Conservatism in human information processing. *Formal representation of human judgment*, 17(51), 1968.

Eisenberger, R. and Weber, M. Willingness-to-Pay and Willingness-to-Accept for Risky and Ambiguous Lotteries. *Journal of Risk and Uncertainty*, 10: 223–233, 1995.

El-Gamal, M. A. and Grether, D. M. Are People Bayesian? Uncovering Behavioral Strategies. *Journal of American Statistical Association*, 90(432):1137–1145, 1995.

Epstein, L. G. and Schneider, M. Learning Under Ambiguity. *Review of Economic Studies*, (74):1275–1303, 2007.

Epstein, L. G. and Schneider, M. Ambiguity, Information Quality, and Asset Pricing. *The Journal of Finance*, 63(1):197–228, 2008.

Epstein, L. G. and Schneider, M. Ambiguity and Asset Markets. 2010.

Epstein, L. G. and Wang, T. Intertemporal Asset Pricing under Knightian Asset Pricing Theory. *Econometrica*, 62(2):283–322, 1994.

Esponda, I. and Vespa, E. Hypothetical thinking and information extraction in the laboratory. *American Economic Journal: Microeconomics*, 2014.

Esponda, I. and Vespa, E. Contingent Preferences and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory. 2016.

Evans, J. S. B. T. *Hypothetical Thinking: Dual Processes in Reasoning and Judgement*. 2007. ISBN 1135419531.

Eyster, E. and Rabin, M. Cursed equilibrium. *Econometrica*, 73(5):1623–1672, 2005.

Eyster, E., Rabin, M., and Vayanos, D. Financial Markets where Traders Neglect the Informational Content of Prices. *Working Paper*, 2015.

Fama, E. F. American Finance Association Efficient Capital Markets A Review of Theory and Empirical Work. *The Journal of Finance*, 25(2):383–417, 1970.

Fehr-Duda, H. and Epper, T. Probability and Risk: Foundations and Economic Implications of Probability-Dependent Risk Preferences. *Annual Review of Economics*, 4(1):567–593, 2012.

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

Fishburn, P. C. Subjective expected utility: A review of normative theories. *Theory and Decision*, 13(2):139–199, 1981.

Forsythe, R. and Lundholm, R. Information Aggregation in an Experimental Market. *Econometrica*, 58(2):309–347, 1990.

Gigerenzer, G. and Hoffrage, U. How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, 102(4):684–704, 1995.

Gilboa, I. and Schmeidler, D. Maxmin Expected Utility with Non-unique Prior. *Journal of Mathematical Economics*, 18:141–153, 1989.

Gilboa, I. and Schmeidler, D. Updating Ambiguous Beliefs. *Journal of Economic Theory*, 59:33–49, 1993.

Gillen, B., Snowberg, E., and Yariv, L. Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort. 2015.

Glosten, L. R. and Milgrom, P. R. R. glosten. *Journal of Financial Economics*, 14:71–100, 1985.

Gonzalez, R. and Wu, G. On the Shape of the Probability Weighting Function. *Cognitive Psychology*, 38:129–166, 1999.

Greene, W. H. *Econometric analysis*. 2008. ISBN 9780131395381.

Greiner, B. An Online Recruitment System for Economic Experiments. In Kremer, K. and Macho, V., editors, *Forschung und wissenschaftliches Rechnen, GWDG Bericht 63*, pages 79–93. Gesellschaft für wissenschaftliche Datenverarbeitung, 2004.

Grether, D. M. Bayes Rule as a Descriptive Model: The Representativeness Heuristic. *The Quarterly Journal of Economics*, 95(3):537–557, 1980.

Grether, D. M. Testing Bayes rule and the representativeness heuristic: Some experimental evidence. *Journal of Economic Behavior and Organization*, 17: 31–57, 1992.

Grossman, S. On the Efficiency of Competitive Stock Markets Where Trades Have Diverse Information. *The Journal of Finance*, 31(2), 1976.

Grossman, S. Further results on the informational efficiency of competitive stock markets. *Journal of Economic Theory*, 18(1):81–101, 1978.

Grossman, S. and Stiglitz, J. E. On the Impossibility of Informationally Efficient Markets. *American Economic Review*, 70(3):393–408, 1980.

Guidolin, M. and Rinaldi, F. A simple model of trading and pricing risky assets under ambiguity: any lessons for policy-makers? *Applied Financial Economics*, 20(1-2):105–135, 2010.

Habib, S., Friedman, D., Crockett, S., and James, D. Payoff and presentation modulation of elicited risk preferences in MPLs. *Journal of the Economic Science Association*, pages 1–12, 2017.

Halevy, Y. Ellsberg revisited: An experimental study. *Econometrica*, 75(2): 503–536, 2007.

Haltiwanger, J. and Waldman, M. Limited Rationality and Strategic Complements: The Implications for Macroeconomics. *The Quarterly Journal of Economics*, 104(3):463–483, 1989.

Haltiwanger, O. and Waldman, M. Rational Expectations and the Limits of Rationality: An Analysis of Heterogeneity. *The American Economic Review*, 75(3):326–340, 1985.

Hanany, E. and Klibanoff, P. Updating preferences with multiple priors. *Theoretical Economics*, 2:261–298, 2007.

Heath, C. and Tversky, A. Preferences and Beliefs: Ambiguity and Competence in Choice under Uncertainty. *Journal of Risk and Uncertainty*, 4(5):5–28, 1991.

Hellwig, M. F. On the Aggregation of Information in Competitive Markets. *Journal of Economic Theory*, 498:477–498, 1980.

Hey, J. and Orme, C. Investigating Generalizations of Expected Utility Theory Using Experimental Data. *Econometrica*, 62(6):1291–1326, 1994.

Hey, J. D. Does Repetition Improve Consistency? *Experimental Economics*, 4: 5–54, 2001.

Hong, H. and Stein, J. C. A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. *The Journal of Finance*, 54(6): 2143–2184, 1999.

Ivanov, A., Levin, D., and Niederle, M. Can Relaxation of Beliefs Rationalize the Winner's Curse?: An Experimental Study. *Econometrica*, 78(4):1435–1452, 2010.

Jaffray, J.-Y. Linear utility theory for belief functions. *Operations Research Letters*, 8(April):107–112, 1989.

Kagel, J. H. and Levin, D. The winner's curse and public information. *American Economic Review*, 76:894–920, 1986.

Kagel, J. H. and Levin, D. *Common value auctions and the winner's curse.* Princeton University Press, 2009.

Kagel, J. H., Levin, D., Battalio, R. C., and Meyer, D. J. First-price common value auctions: bidder behavior and the ”winner's curse”. *Economic Inquiry*, 27(2):241–258, 1989.

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

Kahneman, D. and Tversky, A. Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47(2):263–292, 1979.

Karpoff, J. M. The Relation Between Price Changes and Trading Volume: A Survey. *The Journal of Financial and Quantitative Analysis*, 22(1):109–126, 1987.

Klibanoff, P., Marinacci, M., and Mukerji, S. A Smooth Model of Decision Making under Ambiguity. *Econometrica*, 73(6):1849–1892, 2005.

Klibanoff, P., Marinacci, M., and Mukerji, S. Recursive smooth ambiguity preferences. *Journal of Economic Theory*, 144(3):930–976, 2009.

Koch, C. and Penczynski, S. P. The Winner's Curse: Conditional Reasoning & Belief Formation. 2014.

Kothari, S. P., Shu, S., and Wysocki, P. Do Managers Withhold Bad News? Do Managers Withhold Bad News? *Journal of Accounting Research*, 47(1): 241–276, 2009.

Kyle, A. S. Continuous Auctions and Insider Trading. *Econometrica*, 53(6): 1315–1335, 1985.

Lattimore, P. K., Baker, J. R., and Witte, A. D. The influence of probability on risky choice. A parametric examination. *Journal of Economic Behavior and Organization*, 17:377–400, 1992.

Levin, D., Kagel, J. H., and Richard, J.-F. Revenue Effects and Information Processing in English Common Value Auctions. *The American Economic Review*, 86(3):442–460, 1996.

Lévy-Garboua, L., Maafi, H., Masclet, D., and Terracol, A. Risk aversion and framing effects. *Experimental Economics*, 15(1):128–144, 2012.

Li, S. Obviously Strategy-Proof Mechanisms. 2016.

Loomes, G. and Sugden, R. A Rationale for Preference Reversal. *The American Economic Review*, 73(3):428–432, 1983.

Loomes, G. and Sugden, R. Disappointment and Dynamic Consistence in Choice under Uncertainty. *Review of Economic Studies*, 53(2):271–282, 1986.

Loomes, G. and Sugden, R. Some Implications of a More General-Form of Regret Theory. *Journal of Economic Theory*, 41(2):270–287, 1987.

Loomes, G. and Sugden, R. Testing Different Stochastic Specifications of Risky Choice. *Economica*, 65(260):581–598, 1998.

Madhavan, A. Trading Mechanisms in Securities Markets. *The Journal of Finance*, 47(2):607–641, 1992.

Malinova, K. and Park, A. Liquidity, volume and price efficiency: The impact of order vs. quote driven trading. *Journal of Financial Markets*, 16(1):104–126, 2013.

Miller, G. S. Earnings Performance and Discretionary Disclosure. *Journal of Accounting Research*, 40(1):173–204, 2002.

Mosteller, F. and Nogee, P. An Experimental Measurement of Utility. *Journal of Political Economy*, 59(5):371–404, 1951.

Mukerji, S. and Tallon, J.-M. Ambiguity Aversion and Incompleteness of Financial Markets. *Review of Economic Studies*, (68):883–904, 2001.

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

Ngangoue, K. Trading under ambiguity and the effect of learning. 2016.

Odean, T. Volume, Volatility, Price, and Profit When All Traders Are Above Average. *The Journal of Finance*, LIII(6), 1998.

Olshavsky, R. W. Task complexity and contingent processing in decision making: A replication and extension. *Organizational Behavior and Human Performance*, 24(3):300–316, 1979.

Pagano, M. and Roell, A. Transparency and Liquidity: A Comparison of Auction and Dealer Markets with Informed Trading. *Journal of Finance*, 51(2): 579–611, 1996.

Pires, C. P. A rule for updating ambiguous beliefs. *Theory and Decision*, 53: 137–152, 2002.

Plott, C. R. and Sunder, S. Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets. *Econometrica*, 56(5): 1085–1118, 1988.

Pouget, S. Financial market design and bounded rationality: An experiment. *Journal of Financial Markets*, 10(3):287–317, 2007.

Prelec, D. M. S. S. o. M. The Probability Weighting Function. *Econometrica*, 66(3):497–527, 1998.

Quiggin, J. A theory of anticipated utility, 1982.

Radner, R. Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices. *Econometrica*, 47(3):655–678, 1979.

Ross, B. H. and Murphy, G. L. Category-Based Predictions : Influence of Uncertainty and Feature Associations. *Journal of experimental psychology. Learning, memory, and cognition*, 22(3):736–753, 1996.

Sarin, R. K. and Weber, M. Effects of Ambiguity in Market Experiments. *Management Science*, 39(5):602–615, 1993.

Schnitzlein, C. R. Call and Continuous Trading Mechanisms Under Asymmetric Information: An Experimental Investigation. *The Journal of Finance*, 51(2): 613–636, 1996.

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

Slovic, P. and Lichtenstein, S. Comparison of Bayesian and regression approaches to the study of information processing in judgment. *Organizational Behavior and Human Performance*, 6(6):649–744, 1971.

Starmer, C. and Sugden, R. Violations of the independence axiom in common ratio problems: An experimental test of some competing hypotheses. *Annals of Operations Research*, 19(May 1987):79–102, 1989.

Theissen, E. Market structure, informational efficiency and liquidity: An experimental comparison of auction and dealer markets. *Journal of Financial Markets*, 3(4):333–363, nov 2000.

Tversky, A. and Kahneman, D. Advances in Prospect Theory: Cumulative Representation of Uncertainty. *Source: Journal of Risk and Uncertainty Journal of Risk and Uncertainty*, 5(5), 1992.

Ui, T. The Ambiguity Premium vs. the Risk Premium under Limited Market Participation. *Review of Finance*, 15:245–275, 2011.

Vieider, F. M., Lefebvre, M., Bouchouicha, R., Chmura, T., Hakimov, R., Krawczyk, M., and Martinsson, P. Common components of risk and uncertainty attitudes across contexts and domains: Evidence from 30 countries. *Journal of the European Economic Association*, 13(3):421–452, 2015.

Yates, J. F. and Zukowski, L. G. Characterization of ambiguity in decision making. *Behavioral Science*, 21(1):19–25, 1976.

Zeckhauser, R. Investing in the Unknown and Unknowable. In Diebold, F. X., Doherty, N. A., and Herring, R. J., editors, *The Known, the Unknown, and the Unknowable in Financial Risk Management*, volume 8, pages 1–13. Princeton University Press, 2006.

Ehrenwörtliche Erklärung

Ich habe die vorliegende Dissertation selbstständig und ohne Benutzung anderer als der angegebenen Hilfsmittel angefertigt. Die aus fremden Quellen direkt oder indirekt übernommenen Gedanken sind als solche kenntlich gemacht. Die Dissertation wurde in gleicher oder ähnlicher Form keiner anderen Prüfungsbehörde vorgelegt oder veröffentlicht. Ich bezeuge durch meine Unterschrift, dass meine Angaben über die bei der Abfassung meiner Dissertations benutzten Hilfsmittel, über die mir zuteil gewordene Hilfe sowie über frühere Begutachtungen meiner Dissertation in jeder Hinsicht der Wahrheit entsprechen.

Berlin, den 07. Juni 2017.