Learning under Ambiguity: An Experiment on Gradual Information Processing.*

M. Kathleen Ngangoué[†]

October 2018

Abstract

I investigate the effect of ambiguity on subjects' willingness to trade under different information conditions. The results confirm the prediction of a wide set of theoretical models, that ambiguity aversion reduces willingness to trade in incomplete markets. Participants choose significantly wider bid-ask spreads when return distributions are ambiguous rather than objectively known. This effect also persists when subjects learn probabilities progressively. However, belief updating generates more-extreme quotes that are consistent with a particular updating rule–conditional smooth preferences. These findings highlight the role of gradual information release for belief confidence and under- and overreaction in ambiguous markets.

JEL-Classification: G02, G11, G14, C91, D81, D82
Keywords: learning, ambiguity, market participation, experiment, maximum likelihood updating, full Bayesian updating, conditional smooth preferences.

*A previous version of this paper circulated under the title: "Trading under Ambiguity and the Effect of Learning".

[†]New York University. Email:kn44@nyu.edu

1 Introduction

How do subjects update beliefs under ambiguity? Under risk, that is when the distribution of the states of nature is known, Bayes Rule is the conventional benchmark. In ambiguous environments, however, a multitude of updating rules is conceivable (Jaffray, 1989; Gilboa and Schmeidler, 1993; Epstein and Le Breton, 1993; Pires, 2002; Maccheroni et al., 2006; Epstein and Schneider, 2007; Hanany and Klibanoff, 2007; Hanany et al., 2009; Ghirardato et al., 2008; Klibanoff et al., 2009; Siniscalchi, 2011).

Experimentalists have compared decisions under ambiguity to decisions under risk in static decision settings (see Camerer and Weber, 1992 for a survey). Attitudes toward ambiguity are heterogeneous but the extensive evidence in Ellsberg-type experiments shows that a substantial share of decision-makers dislike ambiguity. In light of these findings, a theoretical literature discusses potential effects of ambiguity aversion on financial decision-making (e.g. Cao et al., 2005; Easley and O'Hara, 2010; Ui, 2011). The ambiguity-averse decision-maker shuns ambiguous settings. In complete markets, for instance, he may reallocate his portfolio to fully insure against ambiguous states. In incomplete markets, reducing his participation in ambiguous markets might be the sole way to avoid ambiguous trades.

Alternatively, in the urge of resolving uncertainty, the ambiguity-averse decision-maker may seek further information. The extent to which ambiguity attitudes are robust to incoming information then depends on the way subjects update beliefs and, as a result, determines whether ambiguity effects persist in the face of market feedback.

Two main paradigms of Bayesian updating under ambiguity –*full Bayesian updating* (henceforth FBU, Jaffray, 1989; Pires, 2002) and *maximum likelihood updating* (henceforth MLU, Gilboa and Schmeidler, 1993)–may result in very different behavior.[1] With FBU, subjects update a set of priors, prior by prior, and retain ambiguity in their posterior beliefs. The relevant beliefs in the resulting set of posteriors are eventually determined by ambiguity preferences. With MLU, on the other hand, subjects con-

[1] Epstein and Schneider (2003) define the condition of rectangularity under which FBU and MLU make identical predictions.

sider a subset of priors that maximizes the *ex-ante* probability of receiving the information. Additional information leads an agent to discard unlikely priors and to perceive substantially less ambiguity. Eventually, he will conceive a single posterior belief and will not perceive any ambiguity at all. In this way, the arrival of information may generate a singleton posterior and eliminate any incentives to avoid ambiguous settings. A third updating rule–*conditional smooth preferences* (CSP)–requires more structure on the preference form but makes predictions in between. Smooth preferences imply that subjects hold beliefs over the set of possible priors. With the arrival of additional information, agents do not only update the support of possible probabilities but also their beliefs over these probability values. Agents may continue to perceive some ambiguity in posteriors but beliefs over the set of posteriors may emphasize some posteriors more than others and, hence, reflect more confidence.

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. A 2x2 design allows me to compare 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 in which information about return distributions is released at once and those in which information is processed sequentially. The design is implemented with two 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. This treatment serves as benchmark to identify subjects' general attitude toward ambiguity in a static framework. Treatment "Learning" (**L**) examines ambiguity effects when investors receive information gradually over time and engage in belief updating.

To understand the impact of learning, the experiment studies investment decisions under ambiguity across two information conditions: one in which investors base their decisions on given probabilities; and another in which investors receive additional information before investing.

The experimental design deviates from standard approaches of measuring ambiguity attitudes with pairwise choices. In a setting that is ubiquitous in financial markets, subjects submit a bid (i.e. their willingness-to-pay) and an ask quote (i.e. their willingness-to-accept a short-sell) 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 then observe an additional signal. The bid-ask spread is then used to compare attitudes toward risk and ambiguity, with and without belief updating.

The result adds to the evidence of ambiguity aversion found in a multitude of Ellsberg experiments (i.e. Chow and Sarin, 2002; Halevy, 2007; and Camerer and Weber, 1992 for a review of the literature). Consistent with ambiguity aversion, participants express a lower willingness to trade by choosing significantly wider bid-ask spreads when returns have ambiguous distributions. 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 and, albeit not surprising, are a necessary benchmark for the comparison with decisions based on updated beliefs.

The results in Treatment L cast doubt on both belief updating theories, FBU and MLU. Ambiguity aversion remains robust to belief updating, showing that subjects were not predominantly MLU agents. MLU predicts small to zero spreads but subjects chose the same average spread when the same ambiguous distribution was learned progressively. The evidence in favor of FBU, too, is limited: although spreads for ambiguous assets with and without learning do not significantly differ, the level of quotes do. Bids and asks are significantly lower (higher) after the arrival of a low (high) signal. Hence, when learning occurs, subjects displayed similar spreads but chose more-extreme quotes. This result is at odds with FBU.

Chosen quotes are rather consistent with updating second-order beliefs about ambiguous probabilities. CSP allows subjects to retain a spread but generates, on average, more-extreme quotes whenever new information is consistent with the expected prior. A bulk of 36.87% decisions for ambiguous prospects was centered around Bayesian updates of the mid-prior. The remainder of quotes disclosed heterogeneity in the way of updating ambiguous beliefs: One noticeable group was insensitive to additional information and refrained from trading; another group behaved like MLU agents by updating extremely and choosing 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. Gradual information processing may nevertheless mitigate ambiguity effects by affecting subjects' confidence in final beliefs. That is, Bayesian updating of recursive preferences may spawn more aggressive bidding despite ambiguity-averse preferences.

Subjects' more extreme reactions to gradual information release have also direct implications for discretionary disclosure policy in financial markets. Miller (2002) and Kothari et al. (2009) find evidence for an asymmetric disclosure of good and bad news: while managers disclose good news immediately, they accumulate bad news before releasing it. The experimental findings indicate that asymmetric disclosure has effects beyond that of supporting managers' careers: it may dampen negative but foster positive stock price reactions.

This paper relates to two strands of research. One strand examines the effect of ambiguity on investment decisions in a static setting. This paper's theoretical predictions is based on Dow and Werlang (1992) that uses ambiguity aversion in form of Choquet expected utility (CEU) to model discontinuity in investors' willingness to trade. Besides CEU, other models of ambiguity aversion (e.g., maxmin expected utility (MEU), α-maxmin expected utility (α-MEU)) also 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, while the ambiguity-averse buyer displays a lower willingness to pay.

Two other experimental studies analyze the effects of ambiguity on financial decisions. Ahn et al.'s (2014) individual-decision experiment confirms the heterogeneity in ambiguity attitudes, providing evidence for subjective expected utility (SEU), ambiguity aversion, and for pessimism. Bossaerts et al. (2010) show in their market experiment that heterogeneity in ambiguity attitudes affects not only portfolio choices but also asset prices. The design in the present experiment deviates from stantard approaches of measuring ambiguity attitudes. Here, the design 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 that bids and the resulting market prices for ambiguous assets are consistently lower in sealed-bid and oral double auctions, although ambiguous and unambiguous assets have identical expected payoffs.[2] As they conclude, subjects are less willing to pay for ambiguous assets that they apparently consider 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 me to test the prediction made in Dow and Werlang (1992) under varying conditions.

Another strand of the literature discusses belief updating under ambiguity. The findings in this experiment support the conjecture in Epstein and Schneider (2007) that information affects the degree of confidence, which in turn may differently affect investment and stock market participation. Cohen et al. (2000) use a dynamic extension of the Ellsberg experiment to differentiate between FBU and MLU behavior. They, too, find heterogeneity in updating behavior. The behavior of a non-negligible number of subjects is consistent with MLU but FBU seems to be the more predomi-

[2]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.

nant 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 also CSP. In De Filippis et al.'s (2016) experiment with both social learning and private signals, subjects' updated beliefs are more consistent with *likelihood ratio test updating*, a generalization of MLU. One 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 presents the stylized decision model and the theoretical predictions. Section 3 describes the implementation of the experiment. The results are presented in Section 4, and Section 5 discusses their implications and concludes.

2 A stylized decision problem

2.1 Investing in ambiguous versus risky prospects

Consider a simple investment opportunity in a market with two states and one asset. The asset has an uncertain value $V \in \{V_L, V_H\}$. The probability for the high-value state corresponds to $Pr(V = V_H) =: \pi$.

The agent is endowed with cash W_0 and has the opportunity to invest in *a single unit* of the asset. He 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 may 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 of buying, as the random price p is more likely to fall below it. Note that the investor will always trade whenever the bid equals the ask. His wealth at the end of the period is $W_1 = W_0 + (V - p)X$.

Denote Π^* as the agent's subjective set of beliefs about π, the probability for the high-value state. We discuss in the following the optimal investment strategy under risk (for EU agents) and ambiguity (for MEU and recursive expected utility (REU) agents with kinked and smooth indifference curves, respectively).

2.1.1 Expected Utility

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 Appendix B.1).

2.1.2 MEU as a model with kinked indifference curves.

Optimal values of bid and ask may change when π is ambiguous. If the agent considers a range 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 quotes. In general, models with kinked indifference curves (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 buying and one for selling. The ambiguity-averse seller short-sells at higher prices, while 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. The present argumentation follows Dow and Werlang (1992) but uses the intuitive model 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 for the two actions of buying and selling. 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 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 (Dow and Werlang, 1992; see Figure 2).[3]

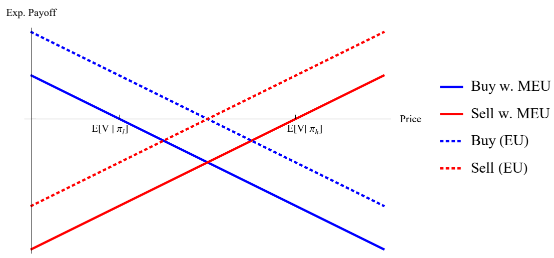

Figure 1: Expected payoff of a buy and a sell as a function of the price for risk-neutral EU (dashed lines) and MEU (solid lines) agents.

[3]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 of hedging the ambiguous asset with the risky one decreases the range of non-participation but does not fully eliminate it (Epstein and Schneider, 2010).

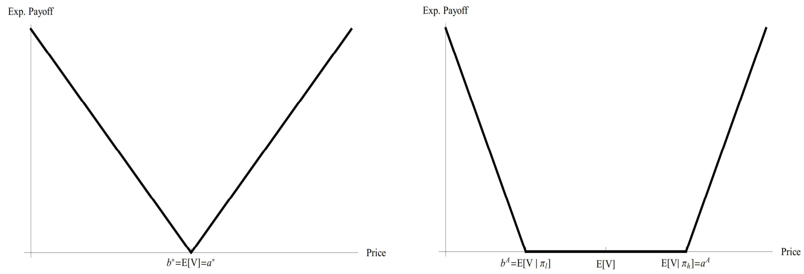

(a) EU **(b)** MEU

Figure 2: Expected payoff with optimal strategy of risk-neutral (a) EU and (b) MEU agents.

2.1.3 REU as a model with smooth indifference curves.

Two of three updating rules discussed in Section 2.2., FBU and MLU, can be applied to preferences with both kinked and smooth indifference curves. The third one, CSP, requires preferences to be smooth. Models of smooth ambiguity preferences, also referred to as *recursive expected utility* (REU) are discussed i.a. in Klibanoff et al. (2005) (henceforth KMM).

An REU agent who considers a set of priors assigns to each single value a probability that it is the true objective probability. In the following, I refer to these beliefs over probabilities as second-order beliefs. Ambiguity is perceived when these second-order beliefs are non-degenerate. Following the model of smooth preferences in KMM, 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:

$$\int_{\pi_l}^{\pi_h} \phi\Big(\mathbb{E}_\pi U(\cdot)\Big)\psi(\pi)d\pi, \tag{1}$$

where $\psi(\pi)$ represents the subjective belief density 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.

As 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 $\phi(\cdot)$.

Section B.2 in the Appendix shows that ambiguity-averse smooth preference produce a wider spread than the spread chosen under risk under the assumption that second-order beliefs are centered around 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.[4] Hence, non-degenerate second-order beliefs will induce the ambiguity-averse subject to choose a wider spread than he would have chosen at the expected prior.

Given risk neutrality, models of ambiguity aversion predict wider spreads for ambiguous than for unambiguous prospects. Predictions under risk aversion, however, depend on the preference model and can vary widely. For preferences with kinked indifference curves, predictions under risk aversion depend on the shape of the indifference curves. For smooth ambiguity preferences such as those studied in KMM, the spread converges to the spread of an expected utility maximizer when ambiguity aversion converges to neutrality. The main objective of the experiment is not to identify kinked versus 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 Introducing belief updating

Consider, now, an environment in which the agent receives an informative signal prior to investing. The signal $s \in \{\vartheta_L, \vartheta_H\}$ is binary, symmetric and

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

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, MEU and REU agents. The difference in predictions also holds under risk aversion.

2.2.1 Bayesian updating

The risk-neutral EU 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, he adjusts the quotes to information but holds a zero spread before and after information. A risk-averse EU agent holds the same non-zero spread for the same belief value, regardless of final beliefs being exogenously given or endogenously updated.

In contrast, if the prior is ambiguous, optimal quotes depend on the way that the agent updates ambiguous beliefs. The literature has proposed various updating rules (Jaffray, 1989; Gilboa and Schmeidler, 1993; Epstein and Le Breton, 1993; Pires, 2002; Maccheroni et al., 2006; Epstein and Schneider, 2007; Hanany and Klibanoff, 2007; Hanany et al., 2009; Ghirardato et al., 2008; Klibanoff et al., 2009; Siniscalchi, 2011). Here, I review *full Bayesian updating* and *maximum likelihood updating*–two main concepts that do not require any specific preference model. Moreover, this two paradigms make maximum opposite predictions with respect to the spread. I also discuss *conditional smooth preferences*–an updating rule that imposes more structure on ambiguity preferences but makes predictions in between.

2.2.2 Full Bayesian updating

Agents with multiple priors apply FBU when they update prior by prior to end up with a set of posteriors. When 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 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 then bids $b = E[V|s = \vartheta_H, \mu_l]$. Analogously, his ask corresponds to $a = E[V|s = \vartheta_H, \mu_h]$, with $b < E[V|s = \vartheta_H, \frac{\mu_l+\mu_h}{2}] < a$.

The MEU agent chooses a non-zero spread both before and after the updating. Its value depends on Π^*, the set of probabilities that he considers possible.

FBU, which does not require second-order beliefs, is a relevant benchmark because the experimental design does not explicitly encourage subjects to conceive second-order beliefs. Moreover, it does not induce a specific shape of second-order beliefs. For completeness, I discuss in the following how an REU agent would apply CSP assuming symmetric second-order beliefs.

2.2.3 Conditional smooth preferences

Assume that second-order beliefs induce an expected prior that equals the midpoint of Π^*, the set of prior values. Albeit simplistic and not induced, this assumption is not unreasonable. For instance, under the principle of insufficient reasons, subjects assign equal probabilities to each of the possible prior values as long as they have no reasons to do differently.

For a simple illustration, assume, in particular, that subjects have uniform second-order beliefs. Upon learning the range of possible priors, they deem every single value within this range equally likely to be the true objective probability. Consider, for instance, uniform second-order beliefs when the prior $\mu \in [0.15, 0.85]$. The arrival of a high signal ($s = \vartheta_H$) affects primarily second-order beliefs: Within the support that is updated upon a high signal, high values become more likely to be the true objective probability than low values. The updating of second-order beliefs will in turn have two implications. First, the high signal shifts the support of possible probabilities to higher values. Second, the asymmetry in second-order beliefs assigns more weight to higher probability values. Figure 3b depicts final second-order beliefs after a low and a high signal.

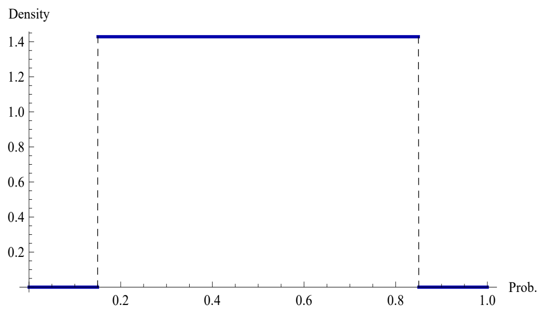

(a) Prior $\mu \in [0.15, 0.85]$

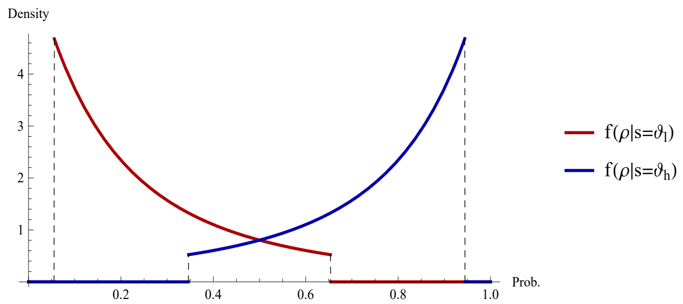

(b) Posterior $\rho \in [.05, .65]$ and $[.35, .95]$

Figure 3: Second-order beliefs over posteriors given uniform second-order beliefs over priors.

Intuitively, processing information has two effects under smooth preferences: incoming information leads to a revision of second-order beliefs that affects not only the support of probabilities but also other moments like the mean belief (see Section B.3 in the Appendix for a formal exposition of these two effects). When incoming information is consistent with prior information, second-order beliefs become more asymmetric and the decision maker holds, on average, more-extreme beliefs. Yet, as long as second-order beliefs remain non-degenerate, the risk-neutral but ambiguity-averse REU agent continues to choose a spread after the updating process.

2.2.4 Maximum likelihood updating

MLU corresponds to an extreme case of CSP in which agents behold only the most likely probabilities. With MLU, the information received pins down the prior that will be updated. The prior that has *ex-ante* the highest probability of generating 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. In that case, signals eliminate any 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, then, satisfies:

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

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

Thus, a fundamental difference between FBU and MLU in this setting is that different factors determine the ranking of states. 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.

2.3 Hypothesis and treatment effect

Both the bid-ask spread and the level of quotes are informative about the updating behavior. However, conclusions about updating behavior can only be drawn if subjects invest differently under risk and under ambiguity–without information update. We therefore first study the effect of ambiguity on chosen quotes.

As discussed in Section 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. Here, the analysis of ambiguity aversion goes beyond any spread increase that can be rationalized 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., $\left(\frac{\pi_l+\pi_h}{2}\right) \neq .50$), a subjective belief of $\Pi^* = .50$ can rationalize a wider spread than the spread chosen at the mid-probability $\left(\frac{\pi_l+\pi_h}{2}\right)$. Note that subjective beliefs fail to rationalize spreads that are wider than any chosen spread at every unambiguous probability $\pi \in [\pi_l, \pi_h]$. The experiment targets evidence in favor of ambiguity aversion that cannot be simultaneously rationalized by subjective expected utility. 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.

Hypothesis 1 *Ambiguous probabilities induce wider bid-ask spreads than unambiguous probabilities:*

$$E[a - b|\pi \in [\pi_l, \pi_h]] > E[a - b|\pi], \quad \forall\, \pi \in [\pi_l, \pi_h]. \tag{2}$$

If subjects are ambiguity-averse, changes in their perception of ambiguity will translate into variation in the spread. In a second step, differences in quotes between the two treatments 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, i.e. rounds with $\Pi^* \in [\pi_l, \pi_h] = [\rho_l^{FBU}, \rho_h^{FBU}]$. Identical spreads and quotes in the two treatments indicate that, on average, subjects perceive the same support of probabilities, which would provide evidence in favor of FBU:

Under FBU, a risk-neutral but ambiguity-averse agent chooses:

i. Identical and strictly positive spreads:

$$0 < E[a-b|\rho \in [\rho_l^{FBU}, \rho_h^{FBU}]] = E[a-b|\pi \in [\pi_l, \pi_h]]$$

ii. Identical quotes:

$$E[a\ |\rho \in [\rho_l^{FBU}, \rho_h^{FBU}]] = E[a\ |\pi \in [\pi_l, \pi_h]]$$
$$E[b\ |\rho \in [\rho_l^{FBU}, \rho_h^{FBU}]] = E[b\ |\pi \in [\pi_l, \pi_h]]$$

with $[\rho_l^{FBU}, \rho_h^{FBU}] = [\pi_l, \pi_h]$.

Conditional smooth preferences (CSP) make predictions in between. A CSP agent considers the same support of posteriors than an FBU agent and, thus, still holds non-degenerate second-order beliefs after the arrival of information. A risk-neutral but ambiguity-averse CSP agent therefore chooses a spread. Average quotes, on the other hand, reflect the expected probability which is more extreme than under FBU. As belief updating generates more asymmetric second-order beliefs, the expected probability becomes more-extreme. These, on average, more-extreme beliefs should be reflected in more-extreme quotes, i.e., quotes that deviate further from 50, the midpoint of the scale.

Under CSP, a risk-neutral but ambiguity-averse agent chooses:

i. Smaller but strictly positive spreads:

$$0 < E[a-b|\rho \in [\rho_l^{FBU}, \rho_h^{FBU}]] < E[a-b|\pi \in [\pi_l, \pi_h]]$$

ii. More-extreme quotes:

$$E[\ |a-50|\ |\rho \in [\rho_l^{FBU}, \rho_h^{FBU}]] > E[\ |a-50|\ |\pi \in [\pi_l, \pi_h]]$$
$$E[\ |b-50|\ |\rho \in [\rho_l^{FBU}, \rho_h^{FBU}]] > E[\ |b-50|\ |\pi \in [\pi_l, \pi_h]]$$

with $[\rho_l^{FBU}, \rho_h^{FBU}] = [\pi_l, \pi_h]$.

Under MLU, beliefs react stronger to information content, which, in turn, eliminates ambiguity. After high (low) signals, beliefs will be higher (lower) than under FBU:

Under MLU, a risk-neutral but ambiguity-averse agent chooses:

i. Zero spreads:

$$0 = E[a - b|\rho \in [\rho_l^{FBU}, \rho_h^{FBU}]] < E[a - b|\pi \in [\pi_l, \pi_h]]$$

ii. Extreme quotes:

$$E[|a - 50|\ |\rho \in [\rho_l^{FBU}, \rho_h^{FBU}]] > E[|a - 50|\ |\pi \in [\pi_l, \pi_h]]$$
$$E[|b - 50|\ |\rho \in [\rho_l^{FBU}, \rho_h^{FBU}]] > E[|b - 50|\ |\pi \in [\pi_l, \pi_h]]$$

with $[\rho_l^{FBU}, \rho_h^{FBU}] = [\pi_l, \pi_h]$.

Thus, comparing the average spread and quotes between treatments NL and L for the same support of marginal and FBU probabilities allows me to differentiate between FBU, CSP or MLU. Table 1 summarizes the predictions in Treatment L for different ambiguity preferences and updating rules.

Table 1: PREDICTED BEHAVIOR OF RISK-NEUTRAL AGENTS WITH DIFFERENT PREFERENCES AND UPDATING RULES

	Risk	Ambiguity			
Preferences	EU	SEU	MEU (or REU)		REU
Updating rule	BU	BU	FBU	MLU	CSP
Spread	0	0	$\|a-b\| = \|\rho_h - \rho_l\| > 0$	0	$0 < \|a-b\| < \|\rho_h - \rho_l\|$
Quotes	$a = b = \rho$	$a = b = \rho$	$a \neq b$	$a = b = \rho \in \{\rho_l, \rho_h\}$	$a \neq b$
			($b = \rho_l \neq a = \rho_h$ under MEU)		$\|\frac{a+b}{2} - 50\| > \|\frac{\rho_l+\rho_h}{2} - 50\|$

Note: For the predictions under ambiguity, we consider the case of ambiguity aversion.

3 Experimental design

3.1 Treatment No Learning (NL)

Treatment **NL** consists of 20 rounds. Subjects start every round with an endowment of cash W_0 and tender both a bid and an ask ($b, a \in [V_L, V_H]$, $b < a$).[5] At the beginning of each round, subjects receive information about the uncertainty of the investment and learn whether or not π is ambiguous. The uncertainty in the asset's value is visualized by displaying "urn A", which 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: the asset takes the value V_L if the value ball is red and 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 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 Figure A1 in Appendix A 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, do not receive any information about how the true composition of urn A is determined when π is ambiguous.

Subjects then quote bid and ask on a second, separate screen.

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.

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

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", which contains 75 pink and 25 green balls. If the value ball is blue, the signal is drawn from "urn H", which contains 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 75% probability.

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 A2 in Appendix A depicts an example of the screen at the second stage.

3.3 Experimental procedures

The computerized experiment was run in the laboratory of Technical University Berlin and WZB Berlin Social Science Center.[6] In total, 67 and 66 students participated in treatments NL and L, respectively. Each treatment was run with three sessions of approximately 22 subjects.

The decision game started once all participants had read the instructions and had responded correctly to a comprehension test. After all subjects completed the decision game, control measures of general attitudes towards risk, uncertainty and ambiguity were elicited (see Appendix Section D).

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

The set of possible probability values was chosen to be parsimonious in order to have enough observations for the comparison between treatments. Each treatment consisted of 14 rounds with unambiguous proba-

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

Table 2: 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].

bilities and six rounds with ambiguous probabilities, or 20 rounds in total. The variation in the unambiguous probabilities π and μ was identical in both treatments NL and L. The ambiguous rounds, on the other hand, differed between the two treatments: in L, the set of priors was fixed to $[.15; .85]$ (see Table 2). There, the variation in beliefs came from the signal's value that implied 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 2.3, the two sets of probabilities, $[.05; .65]$ and $[.35; .95]$, in NL were chosen to equal the set of posterior beliefs under FBU in L. This enables me to compare bids and asks for the same dispersion in probabilities, when information on the distribution was provided immediately versus sequentially.

Within each treatment, participants made their decisions in alternating blocks of seven consecutive risky and three 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 was reversed. In addition, subjects played four 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 decision game (Baillon et al., 2015). For this purpose, subjects threw a 20-sided die after the trial rounds but before playing the 20 rounds. Subjects did not see the outcome of the die roll until the end of the game. That is, they were aware that the payoff-relevant round was fixed during the experiment but they learned which round was chosen only after all of their decisions. The instructions as well as the computer screen emphasized accordingly that hedging across rounds makes no sense once the payoff-relevant round is determined.

Earnings consisted of a show-up fee (5 EUR), plus two thirds of earnings in the randomly drawn round in the investment game plus one third of earnings in 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.

4 Results

4.1 Treatment NL

Decisions for risky prospects. Subjects made mostly risk-averse choices: a majority of bid-ask pairs had a non-zero spread. Since the distribution of spreads is highly right-skewed, the analyses focus mainly on quantiles.[7] 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 4a). The spread is asymmetric around the probability, reflecting that increasing the bid (the ask) becomes more (less) risky with an increasing probability (see Figure 4b).[8][9] Overall, subjects chose a median spread of 5 ECU.

Table 3: 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 significantly reduced subjects' willingness to trade. The median bid is shifted downwards, and the median ask increases, leading to significantly

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

[8] 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 of paying 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. However, the asymmetry in risk premia observed in the data is too large to be rationalized by reasonable coefficients of risk aversion.

[9] Subjects are generally more risk-averse in buying than in selling. This finding is puzzling as endowment effects should not play a role in short-selling.

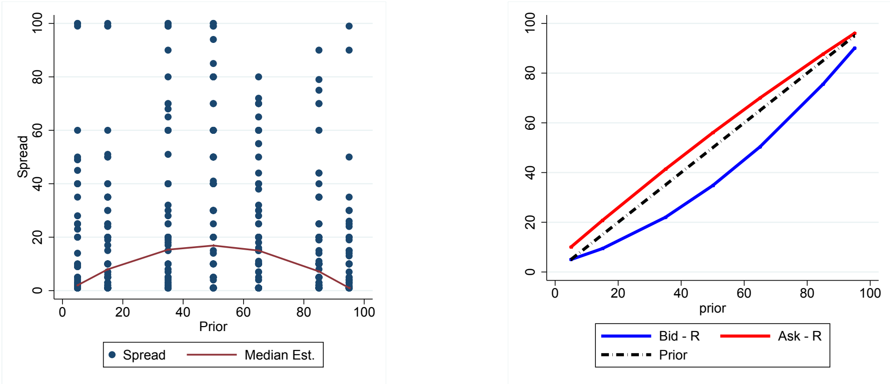

(a) Spreads and median regression

(b) Median bid and ask.

Figure 4: Median spreads and quotes as a function of unambiguous priors.

wider spreads for ambiguous prospects (see Table C3 in Appendix C.1.2). Median spreads for prospects with ambiguous probabilities are three times as high as for unambiguous probabilities (see Table 3). Despite the asymmetry in risk premia for long and short positions, the ambiguity premium is almost symmetric. Subjects exhibited 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, respectively (see Table C2 in Appendix C.1.1). In sum, Hypothesis 1 is confirmed.

Result 1 *Ambiguity in probabilities generates wider spreads.*

As a direct consequence of the design, subjects traded and earned less when the return distribution was ambiguous. Subjects traded risky prospects in 82% of all rounds. Trades fell by 14.8% (12 percentage points) when probabilities were ambiguous. The greatest reduction of 19.3% occurred when the probability was between 15% and 85% (see Table 4).

The decrease in trading activity translated into significantly smaller profits. Subjects earned, on average, 41.98% (p=.0015, two-sample t-test) more in risky rounds than in ambiguous rounds (See Table C1 in Appendix C.1.1).

Table 4: 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.

4.2 Treatment L

4.2.1 Spreads

The general effects of ambiguity on the spread are robust to incoming information. In the aggregate, choices in treatment L were ambiguity-averse. Subjects chose wider spreads for ambiguous than for risky distributions, with increasing difference in the mean in the last ten rounds (see Table 5).

Table 5: 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.

The comparison between the two treatments shows that subjects still reacted to information. Starting with a set of priors $\mu \in [.15, .85]$, full Bayesian inference reduces the interval of probabilities by ten 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 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 (on average 6.29 ECU more, p=0.088, two-sample test).

However, no difference in the aggregate distribution of spreads is observable after controlling for the range of marginal and FBU probabilities (p-value=.92 in Kolmogorov-Smirnov test; see Figures 5a and 5b). Comparing rounds in which marginal probabilities (π) and FBU posteriors (ρ) lie in the same interval [.05; .65] discloses a small difference in the spread: participants in NL chose a median spread of 20, whereas the median spread in L equaled 15. This non-significant difference carries even less weight in the aggregate since the two treatment groups chose identical median spreads of 20 when both π and $\rho(s = \vartheta_h) \in [.35; .95]$.

Figure 5a depicts the distribution of chosen spreads in the ambiguous rounds of treatment NL with $\pi \in [.05, .65]$ or [.35; .95]. Figure 5b refers to the distribution 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).

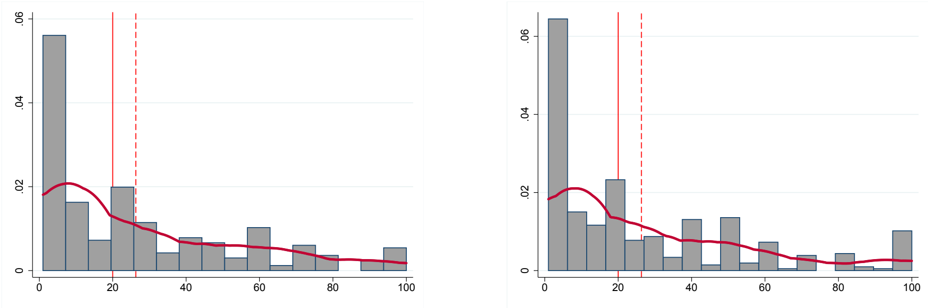

(a) Ambiguous rounds in NL **(b)** Ambiguous rounds in L

Figure 5: Spreads for ambiguous prospects for the same theoretical dispersion in marginal (a) and Bayesian posterior (b) probabilities.

Apparently, subjects did not perceive substantially less ambiguity when the same information was released gradually. Therefore, the data do not lend support to MLU theory. Yet the data are not completely consistent with FBU theory, either: subjects reacted differently to ambiguity in given probabilities than to ambiguity in updated posteriors. Although spreads were, in the aggregate, constant, chosen bids and asks were more extreme after information. The next section analyses the level of quotes as an additional indicator about updating behavior besides the spread.

4.2.2 Quotes

Learning generated more-extreme quotes. Participants in treatment NL chose a median bid and ask of 17.5 and 50 when $\pi \in [.05; .65]$. Participants in treatment L, however, chose a median bid and ask of 10 and 40 for an FBU posterior $\rho \in [.05; .65]$ (significant differences between updated and non-updated quotes at the 5% level each). Analogously, the median bid and ask is 40 and 70.5 in the rounds in which $\pi \in [.35; .95]$ but 50 and 81 in the rounds with a set of FBU posteriors $\rho \in [.35; .95]$ (significant differences at the 1% level each).

The process of updating also introduced heterogeneity in chosen quotes. Yet this heterogeneity did not seem to be driven by a lack of probabilistic sophistication. Section C2 in the Appendix shows that Bayesian inference cannot be rejected in risky rounds. Estimated decision weights reflect an inverse S-shape weighting function. Taking into account the estimated weighting function, updated quotes conformed, on average, with Bayesian inference (see Appendix Figure C1).

As subjects, in the aggregate, updated information correctly in risky rounds, the heterogeneity in ambiguous rounds cannot be reduced to mistakes and noise. It might rather reflect different updating rules. To illustrate the heterogeneity in quotes, the midpoints of bid and ask pairs (henceforth mid-quotes) are depicted in Figures 6a-6d. The top two panels, 6a and 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, 6c and 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 three mass points ($\{0-5; 20-25; 45-50\}, \{50-55; 70-75; 95-100\}$), suggesting three main updating methods.

The cluster analysis in Appendix C.3 illustrates how decisions differed. In sum, a substantial share of quotes (25.75%) matched highly ambiguity-averse investment behavior, which favored non-participation. These subjects centered their bids and asks around the mid-prior 50 and chose wide spreads. Another substantial share (21.72%) were consistent with MLU: quotes were extreme and spreads minimal. The majority of decisions (42.17%) reflected updated but less extreme quotes. Yet these Bayesian quotes did not reflect FBU posteriors. Under FBU, participants in treatments NL and L should have considered the same support of probabilities and, therefore, made similar decisions. Bid-ask pairs in treatment L should have resembled the ones in NL and should have been similarly centered around midpoints of the sets of probabilities, which, here, were $\{35, 65\}$. However, bid-ask quotes based on incoming information reflected more extreme beliefs than the ones in treatment NL. Controlling for the range in marginal and FBU probabilities, 36.19% of bid-ask pairs in treatment NL encompassed the value 50 versus 29.54% in treatment L (p-value=0.07 in binomial test).

Smooth preferences account for the difference between average quotes in treatments L and NL. Conditional smooth preferences generate more-extreme beliefs than marginal smooth preferences if traders have mean-

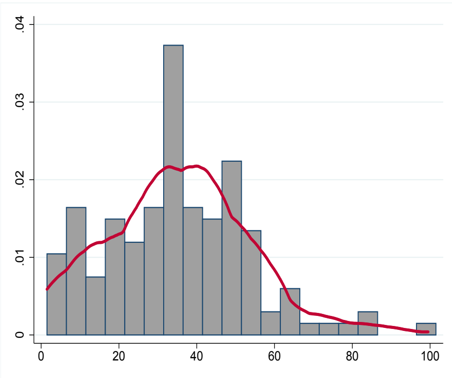

(a) $\pi \in [.05, .65]$

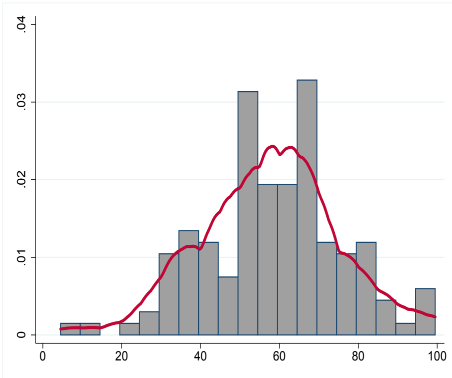

(b) $\pi \in [.35, .95]$

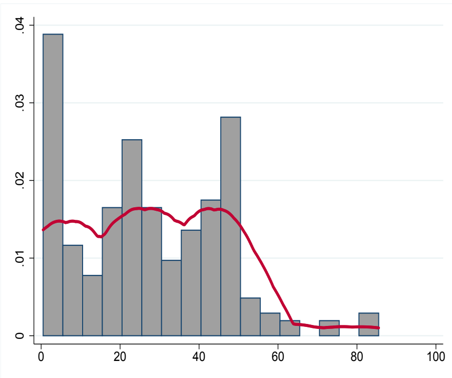

(c) $\rho \in [.05, .65]$

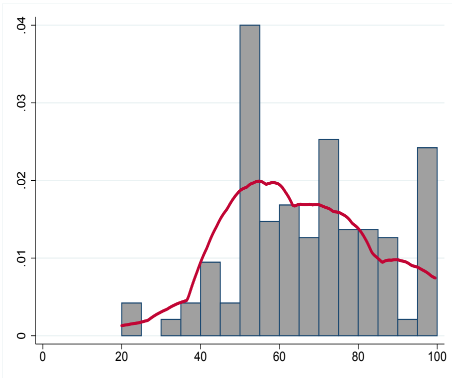

(d) $\rho \in [.35, .95]$

Figure 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.

preserving second-order beliefs. Consequently, gradual information release induces more-extreme quotes compared to an environment in which information is released all at once. Figures 7a and 7b display second-order beliefs with and without learning for the same support of probabibilities. 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.

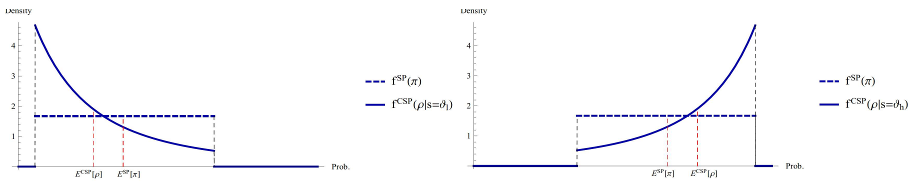

(a) $\pi, \rho \in [.05, .65]$ **(b)** $\pi, \rho \in [.35, .95]$

Figure 7: Marginal and Bayesian second-order beliefs for a low (a) and a high (b) support of probabilities.

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.[10] Average chosen quotes can indeed be rationalized with updating the prior $\pi = 0.5$, the midpoint in the set of priors (see Figure of NLSUR in Appendix). In this context, Bayesian updating of second-order preferences illustrates why quotes are more extreme in treatment L than in NL for the same support of probabilities.

[10]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 adjusted their quotes to the prior in risky rounds. Figure C4 in Appendix C.2.2 shows how mid-quotes increase in the prior for the different signal values. Bids and asks are not heavily centered around 25 or 75.

Intuitively, processing information with smooth preferences entails a revision of second-order beliefs, which, in turn, has several implications: When incoming information is consistent with prior information, second-order beliefs become asymmetric, generating both an updated support of probabilities and more-extreme average first-order beliefs.

5 Conclusion

The evidence of ambiguity aversion found so far in Ellsberg-type experiments extends to other frameworks. The experiment shows that, when 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 more risky (Sarin and Weber, 1993; Epstein and Wang, 1994).

A second main insight from 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, subjects chose more-extreme quotes when they received information in pieces.

In addition, incoming information introduces more heterogeneity into trading behavior. A substantial share of subjects were insensitive to additional information; another non-negligible share adopted extreme beliefs; and the majority of subjects appeared to update second-order preferences in a Bayesian way.

More questions remain to be clarified in future research. First, ambiguity effects may 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 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, one might find it sufficient to not be at an informational disadvantage compared to other traders. Furthermore, the interaction between investors might eliminate any perception of ambiguity, especially if aggressive traders dominate the markets. Indeed the evidence on ambiguity effects in ambigous experimental markets is mixed (Sarin and Weber, 1993; Bossaerts et al., 2010; Kocher and Trautmann, 2013; Corgnet et al., 2013; Füllbrunn et al., 2014) and 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. This is consistent with Epstein & Schneider's (2007) conjecture that increasing confidence through learning fosters investment and stock market participation. This study draws attention to frequent information release as a mechanism to alter subjects' confidence in their final beliefs and avoid or correct frictions in trades.

Acknowlegments

I am deeply grateful to Georg Weizsäcker and Antonio Guarino for their guidance and support. This project also benefited from the helpful comments of Marina Agranov, Pio Baake, Yves Breitmoser, Peter Cramton, Erik Eyster, Dan Friedman, Itzhak Gilboa, Yoram Halevy, Philippe Jehiel, Peter Klibanoff, Matthias Lang, Andrew Schotter and the participants of seminars at several universities. All errors are mine. I thank the Humboldt-Universität zu Berlin and the DIW Berlin for their financial support.

References

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

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

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

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

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

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

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

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

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

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

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

Corgnet, B., P. Kujal, and D. Porter (2013). Reaction to public information in markets: How much does ambiguity matter? *Economic Journal 123*(569), 699–737.

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

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

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

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

Epstein, L. G. and M. Le Breton (1993). Dynamically consistent beliefs must be bayesian.

Epstein, L. G. and M. Schneider (2003). Recursive multiple-priors. *Journal of Economic Theory 113*(1), 1–31.

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

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

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

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

Füllbrunn, S., H. A. Rau, and U. Weitzel (2014). Does ambiguity aversion survive in experimental asset markets? *Journal of Economic Behavior and Organization 107*(PB), 810–826.

Ghirardato, P., F. Maccheroni, and M. Marinacci (2008). *Revealed Ambiguity and Its Consequences: Updating.* Berlin Heidelberg: Springer-Verlag 2008.

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

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

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

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

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

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

Hanany, E., P. Klibanoff, E. Hanany, and P. Klibanoff (2009). Updating Ambiguity Averse Preferences Updating Ambiguity Averse Preferences. *The B.E. Journal of Theoretical Economics 9*(1), 1–51.

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

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

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

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

Kocher, M. G. and S. T. Trautmann (2013). Selection into auctions for risky and ambiguous prospects. *Economic Inquiry 51*(1), 882–895.

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

Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Dynamic variational preferences. *Journal of Economic Theory 128*(1), 4–44.

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

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

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

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

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

Siniscalchi, M. (2011). Dynamic choice under ambiguity. *Theoretical Economics 6*(3), 379–421.

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

Vieider, F. M., M. Lefebvre, R. Bouchouicha, T. Chmura, R. Hakimov, M. Krawczyk, and P. Martinsson (2015). 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.

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

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

For Online Publication (Appendices A-D)

A Screen layout

Figures A1a and A1b depict examples of the composition in urn A when the prior is unambiguous and ambiguous, respectively. The grey balls in the ambiguous urn can be either red or blue.

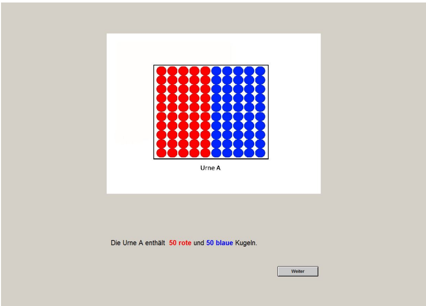

(a) Risky prospect

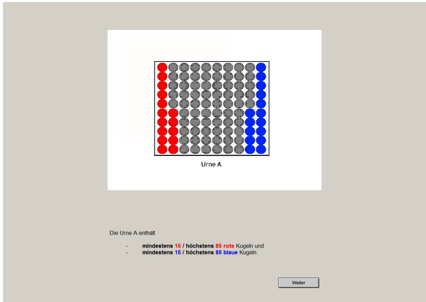

(b) Ambiguous prospect

Figure A1: Examples for visualization of probability distribution with urn A.

In treatment L, them subjects viewed a second decision screen before they chose their quotes (see Figure A2). In the upper left corner, the composition in urn A reminded 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 "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.

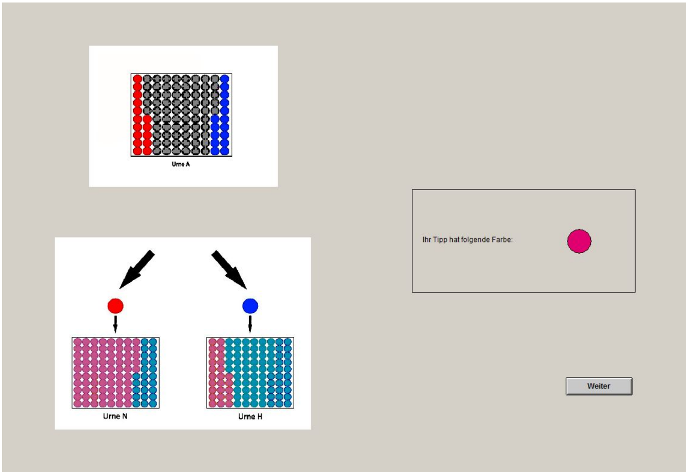

Figure A2: Example for an additional signal at the second stage.

B Mathematical appendix

B.1 Bid-ask spread generated by risk aversion

Risk aversion introduces a spread between the bid and the ask.

Let $b^{RN} = \mathbb{E}(V)$ be the optimal bid under risk neutrality. Assume that risk-averse preferences are 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 investment in 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.2 Bid-ask spread with ambiguity-averse smooth 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:

$$\int_{\pi_l}^{\pi_h} \phi\Big(\mathbb{E}_\pi U(\cdot)\Big)\psi(\pi)d\pi, \tag{3}$$

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.

The following analysis assumes that subjects have second-order beliefs whose mean corresponds to the midpoint in the range of priors.[11] I first show that ambiguity-averse but mean-preserving second-order preferences generate a bid-ask spread. The optimal bid for going long is the certainty equivalent that satisfies:

$$\int_{\pi_l}^{\pi_h} \phi\Big(\mathbb{E}_\pi U(W_0 + V - b)\Big)\psi(\pi)d\pi = \phi(U(W_0)). \tag{4}$$

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

$$\mathbb{E}_\psi \phi(\mathbb{E}_\pi U(W_0 + V - b)) < \phi(\mathbb{E}_\psi \mathbb{E}_\pi U(W_0 + V - b)). \tag{5}$$

Under mean-preserving second-order beliefs, the subjective probability function $\psi(\pi)$ 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 (5) then equals:

$$\phi(\mathbb{E}_\psi \mathbb{E}_\pi U(W_0 + V - b)) = \phi(\mathbb{E}_{\bar{\pi}} U(W_0 + V - b)). \tag{6}$$

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

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 :

$$\phi(\mathbb{E}_{\bar{\pi}}U(W_0 + V - b^R) = \phi(U(W_0)). \tag{7}$$

From equations (4), (5) and (7), it follows that:

$$\phi(\mathbb{E}_{\bar{\pi}}U(W_0 + V - b^R) < \phi(\mathbb{E}_{\bar{\pi}}U(W_0 + V - b)). \tag{8}$$

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 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.

B.3 Conditional smooth preference

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 directly affects 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

$\psi(s,\mu) \neq \psi(\mu)$:

$$\mathbb{E}_{\psi(s=\vartheta_l,\mu)}\phi\Big(\mathbb{E}_{\{s=\vartheta_l,\mu\}}U(\cdot)\Big) < \mathbb{E}_{\psi(\mu)}\phi\Big(\mathbb{E}_{\{s=\vartheta_l,\mu\}}U(\cdot)\Big) \tag{9}$$

$$\mathbb{E}_{\psi(s=\vartheta_h,\mu)}\phi\Big(\mathbb{E}_{\{s=\vartheta_h,\mu\}}U(\cdot)\Big) > \mathbb{E}_{\psi(\mu)}\phi\Big(\mathbb{E}_{\{s=\vartheta_h,\mu\}}U(\cdot)\Big). \tag{10}$$

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 in which information is released all at once.

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.

C Results

C.1 Reactions to ambiguity

C.1.1 Descriptive statistics

Table C1 shows mean profits for risky and ambiguous prospects, across different ranges of probabilities.

Table C1: 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 dummy variable for the ambiguous rounds.

Table C2 shows the 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 C2: 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.

C.1.2 Regression estimates

Table C3 presents the results of the median polynomial regression. The estimates for risky prospects are plotted in the Figures 4a and 4b in Section 4.1.

Table C3: 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.

C.2 Probabilistic sophistication

Updating unambiguous priors. Subjects' general probabilistic sophistication is analyzed with their 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 (NLS-SUR):

$$\begin{cases} b_i = (1 - RP_b) \cdot E[V|\tilde{\tau}] + \epsilon_{i,b} \\ a_i = (1 + RP_s) \cdot E[V|\tilde{\tau}] + \epsilon_{i,s}, \end{cases} \tag{11}$$

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 were 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 by 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 $\tau = \rho$.[12] The coefficient α regulates the curvature of the function. The parameter β determines the inflection point of the curve.

Table C4: 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 reacted to unambiguous marginal probabilities in the same way as

[12] 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 objective probabilities.

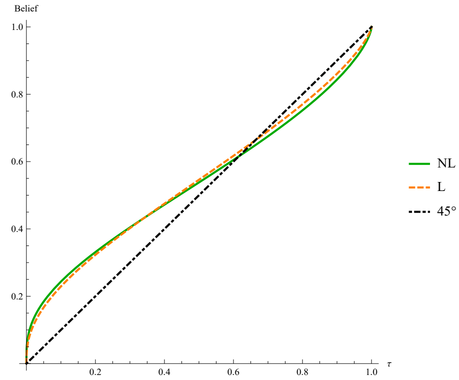

Figure C1: Estimated probability weighting function for unambiguous probabilities in NL & L.

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 ambiguous 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, which 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 C2a and C2b depicts the relation between subjects' estimated beliefs and ambiguous probabilities in NL. This inverse s-shape relation serves as a benchmark for the relation between estimated beliefs and the midpoint of ambiguous posterior probabilities in L. The dashed line in Figure C2a represents the model fit with FBU posteriors. Estimated beliefs are 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, chosen quotes were too extreme to be explained by the range of beliefs under FBU.

The dashed line in Figure C2b depicts the model fit with MLU posteriors. The weighting function is inverse s-shaped but also deviates from the benchmark (solid line). Given an MLU probability, estimated beliefs are not sufficiently extreme to match the benchmark. Decision weights are too close to the belief of 50% to be explained by extreme MLU posteriors. Appendix Section C.2.1 displays the estimates of the NLS-SUR with ambiguous probabilities and the results of a Lagrange-Multiplier test, which show a significant difference between the benchmark model and the model fit under both FBU and MLU probabilities.

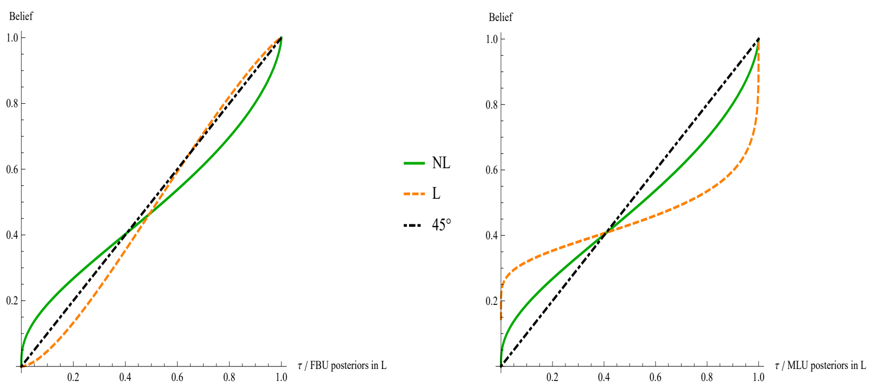

(a) Assuming FBU in L. **(b)** Assuming MLU in L.

Figure C2: Estimated probability weighting functions in ambiguous rounds of 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.

Instead, Bayesian posteriors at the expected prior $\rho(s, \pi = .5)$ fit the relation between trading decisions and probabilities (see Figure C3): 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.

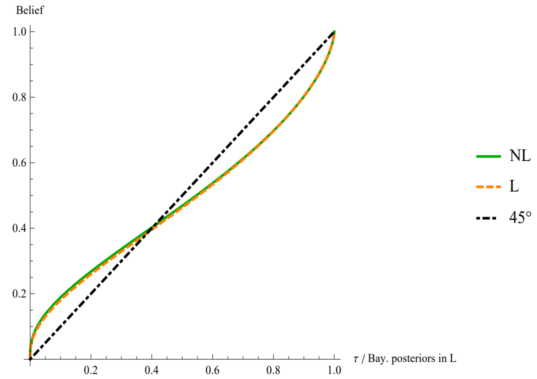

Figure C3: Estimated probability weighting functions in ambiguous rounds of NL & L assuming BU of mid-prior in L.

C.2.1 Results of NLS-SUR

Table C5 shows the coefficient estimates of the NLS-SUR model. Because, by design, there is less variation in the ambiguous probabilities, the estimation is more efficient when assuming symmetric premia in the bid and the ask. The model estimates assuming Bayesian update of recursive preferences (CSP - conditional smooth preferences) 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 C5: Coefficient estimates for probability weighting function and risk premia

	NL	L		
	mid-prior	FBU***	MLU***	CSP
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.

C.2.2 Heterogeneity in updating

Figure C4 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.

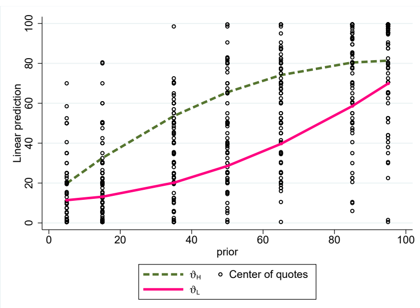

Figure C4: Mid-quotes for unambiguous assets and their mean estimates for the two signals and the group of Bayesian updaters (clusters 3 and 5 in Section E.3).

C.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 eight clusters, yielding the eight different ranges for updated beliefs listed in Table C6.[13]

Table C6: MEDIAN BIDS, ASKS AND SPREADS AND CORRESPONDING STATISTICS FOR EIGHT 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.

In total, 21.72% of the ambiguous decisions belong to clusters 1 and 6 and are consistent with MLU. Quotes in these clusters were close to one extremum and exhibited, on average, the smallest spread of one ECU. The opposite behavior is described in clusters 7 and 8, which represent 25.75% of the bid-ask pairs. These observations exhibited a substantial spread of more than 30 ECU. In approximately one third of these decisions, the chosen spread was wide enough to almost surely implement a no-trade outcome (cluster 8). In cluster 4, 10.35% of the quotes disclosed a small spread with bids and asks around 50%, the midpoint of the set of priors. These quotes

[13]The value of eight clusters finds its justification in the theory, allowing the identification of eight 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); and ambiguity-neutral likelihood-insensitive beliefs (50,50). Robustness checks with more and fewer clusters do not yield better comprehension of the data.

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 amount to 42.17% of bid-ask pairs in 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, fall in clusters 3 or 5, which disclose a median spread of 20 ECU. Figure C5 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.

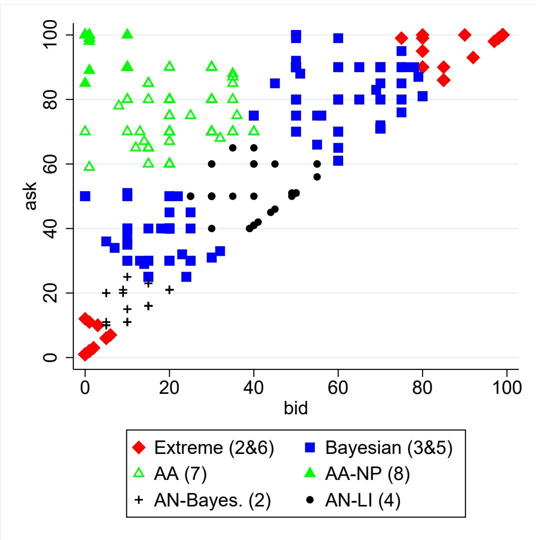

Note to the legend: Cluster categories in parentheses.

Figure C5: Clusters of bid-ask pairs in ambiguous rounds of treatment L

Table C7 lists the results of the same cluster analysis in treatment NL. The analysis yields less extreme clusters of beliefs. Furthermore, the observations are distributed more evenly across the eight clusters, yielding the more symmetric distributions of quotes reflected in Figures 6a and 6b.

Table C7: 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.

D Elicitation of uncertainty preferences

I elicited control measures of preferences to analyze the extent to which subjects' behavior conformed with standard measures of risk and ambiguity attitudes. For a cleaner comparison with the pricing task in the main experiment, I used certainty equivalents as a main measure for risk and ambiguity preferences. Feedback on payoff was provided only after completion of Part 2. All measures were elicited by displaying multiple price lists with increasing numbers from top to bottom. This lack of randomization might bias attitudes in a systematic direction, for instance if subjects have an inclination to choose rows at the top or at the bottom. The results in this section require therefore a cautious interpretation.

D.1 Risk attitudes

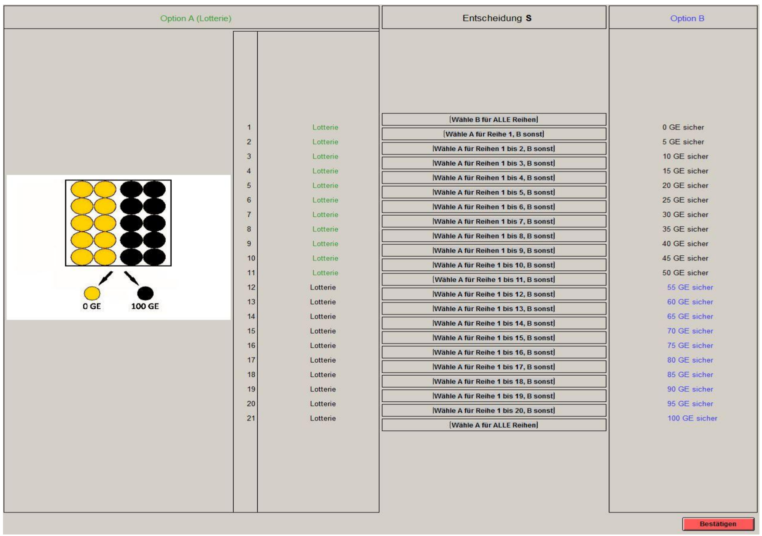

Figure D6: Example of computer interface in Part 2

Risk preferences were elicited with a multiple price list task akin to Abdellaoui et al. (2011) and Gillen et al. (2015). In two replicate measurements, subjects faced 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 corresponded to $(100, 0.5; 0)$ and $(150, 0.5; 0)$, respectively. The lottery was illustrated on the left side of the computer interface, where subjects saw an urn with 10 (15) yellow and 10 (15) black balls in the first (second) measurement. The lottery payed out if a black ball was drawn. The right side of the interface showed a list of sure payoffs in $[0; x]$, with increments of 5 ECU per row. Subjects must then, for each row, make a pairwise choice between the lottery and the sure payoff. Monotonicity was enforced as subjects could only switch once from preferring the lottery to preferring a sure payoff. Figure D6 depicts the computer interface for the first measurement with lottery $(100, 0.5; 0)$.

D.2 Uncertainty attitudes

Uncertainty attitudes were measured in two settings: in one task subjects stated their certainty equivalent for a lottery with unknown probabilities; the other task refers to a standard Two-Urn Ellsberg-Experiment in which subjects chose between a risky and an ambiguous lottery.

D.2.1 Certainty equivalent

Subjects were presented with the same two multiple price list choices as in the elicitation of risk attitudes but lotteries had unknown probabilities. An urn with 20 (30) grey balls was used to illustrate the lottery with unknown probabilities. Note, however, that the elicitation of certainty equivalents for a bet on a black ball does not enable identifying ambiguity aversion. A subject with a pessimistic belief would choose a low certainty equivalent without being necessarily ambiguity averse. Ambiguity aversion requires aversion towards uncertainty for both sides of the bet.

D.2.2 Two-Urn Ellsberg problem

Subjects made choices involving two lotteries with a high prize of either 100 or 150 ECU depending on the urn size. For each lottery, they faced two gambles: A bet on a yellow ball that would pay 100 (150) ECU if a yellow ball was drawn from an urn with 20 (30) balls and a bet on a black ball that would pay 100 (150) ECU if a black ball was drawn from the *same* urn. In a multiple price list, subjects specified their preferences between urn I and urn II. The two urns had a total of 20 (30) balls, in a combination of yellow and/ or black balls. While the composition of yellow and black balls were unknown in urn I, the composition in urn II varied along the list. For a bet on a yellow ball, subjects indicated the minimum amount of yellow balls in urn II, for which they were willing to switch from urn I to urn II. Analogously for a bet on a black ball, they specified the minimum amount of black balls in urn II. In the following, the term "matching probability" refers to the share of balls at which subjects started to prefer the risky lottery.

D.3 Results

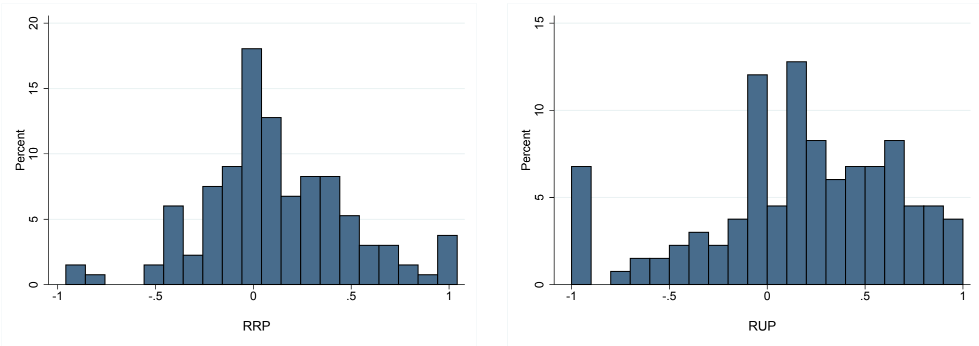

(a) RRP

(b) RUP

Figure D7: Relative risk and uncertainty premia

I define the certainty equivalent (CE) as the midpoint of the two payoffs between which subjects switched from preferring the lottery to the sure payoff. Figures (D7a) and (D7b) depict the distribution of the relative risk and relative uncertainty premia ($RRP = RUP = \frac{E(x)-CE}{E(x)}$). The uncertainty premium is measured relative to a success probability of 50%. On average, 57.9% and 65.4% of all subjects chose a positive risk and uncertainty premium, respectively. A share of 26.3% had neither a positive risk nor uncertainty premium, 15.8% displayed a positive uncertainty premium but no risk premium.

Let γ_w, γ_b be the matching probabilities for bets on yellow and bets on black balls. The degree of ambiguity aversion is measured by $\delta = (\gamma_w + \gamma_b)$.

$$\delta \begin{cases} > 1 & \texttt{ambiguity seeking} \\ = 1 & \texttt{ambiguity neutrality} \\ < 1 & \texttt{ambiguity aversion} \end{cases}$$

.

The elicited ambiguity attitudes are not consistent with the ambiguity aversion reflected in chosen spreads and the, on average, positive uncertainty premia. Only 18.8% of subjects were ambiguity averse but a surprisingly large fraction of 72.93% subjects were ambiguity seeking. This casts some doubts on the elicitation procedure and the measures' robustness. Ambiguity attitudes were elicited last and it is not clear whether this result is due to fatigue, experience, some misunderstanding or framing in the interface (which was not randomized).

In general, chosen spreads were consistent with the elicited measures of risk and uncertainty premia but not with the elicited ambiguity attitudes. Subjects who displayed a positive risk and uncertainty premium chose significantly wider spreads. However, the elicited ambiguity attitudes do not correlate with chosen spreads (correlation coefficient of -0.01).

Table D8: Average spreads for different categories of elicited attitudes

		mean spread	med. spread
RRP	< 0	15.43	4
		(0.80)	
	> 0	21.62***	10
		(0.82)	
RUP	< 0	21.58	13
		(1.49)	
	> 0	31.08***	20
		(1.31)	
δ	< 1	25.16	20
		(2.06)	
	$= 1$	25.03	20
		(3.20)	
	> 1	28.79	20
		(1.22)	

Note: *** denote significant differences in spreads between subjects with positive and nonpositive premia, with a p-value < 0.01.