Learning under Ambiguity:
An Experiment on Gradual Information Processing.

M. Kathleen Ngangoué[a]

June 2020

Abstract

This experiment studies belief updating under ambiguity, using subjects' bid and ask prices for an asset with ambiguous payoff distribution. Bid and ask quotes allow for distinguishing between the two main paradigms of updating under ambiguity–*full Bayesian updating* and *maximum likelihood updating*. We find substantial heterogeneity in updating behavior. The majority of subjects (54%) chose quotes that were in line with *full Bayesian updating*, but another non-negligible group (35%) behaved like *maximum likelihood* updaters.

JEL-Classification: G11, C91, D81
Keywords: learning, ambiguity, experiment, maximum likelihood updating, full Bayesian updating.

[a]New York University, 19 West 4th St., 6th floor, NY NY 10012. 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 where agents may contemplate more than a single distribution, a multitude of updating rules are conceivable (Jaffray, 1992; 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).

In this paper we study belief updating in a context of multiple priors. The two main paradigms of Bayesian updating with multiple priors –*full Bayesian updating* (henceforth FBU, Jaffray, 1992; Pires, 2002) and *maximum likelihood updating* (henceforth MLU, Gilboa and Schmeidler, 1993)–may spawn very different decisions.[1] With FBU, subjects update a set of priors, prior by prior, and retain ambiguity in their posterior beliefs. With MLU, on the other hand, subjects consider a subset of priors that maximizes the *ex-ante* probability of receiving the information. In other words, additional information leads an agent to discard unlikely priors and to perceive substantially less ambiguity. Eventually, the arrival of information may prompt the agent to conceive a unique probability measure, and thus eliminate her incentives to avoid ambiguity. This paper offers an experimental design that makes use of this contrast between FBU and MLU to identify decisions that are consistent with one or the other.

Updating behavior under ambiguity is still little understood (see the discussion on related literature) and more empirical evidence is needed, not only because ambiguity is ubiquitous in various decision problems but, in particular, because it is precisely ambiguity that might induce decision makers to seek further information. Whether or not information processing disproportionally reduces ambiguity is an important question given the broad evidence of ambiguity effects on decisions. So far, a large empirical literature has focused on the effects of ambiguity on decision-making while abstracting from belief updating. In general, attitudes toward ambiguity are heterogeneous but the extensive evidence in Ellsberg-type experiments shows that a substantial share of decision-makers dislike ambiguity.[2] In light of these findings, a theoretical literature discusses potential effects of ambiguity aversion on financial decision-making (e.g. Cao

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

[2]See also Machina and Siniscalchi (2014); Marinacci (2015); Gilboa and Marinacci (2016) for excellent reviews of the theoretical literature and their applications.

et al., 2005; Easley and O'Hara, 2010; Ui, 2011). The ambiguity-averse decision maker shuns ambiguous settings; in financial markets, for instance, through portfolio reallocation or limited participation. Yet, the importance of ambiguity effects in real decision contexts clearly depends on the extent to which ambiguity persists in a dynamic setting where agents are likely to engage in information acquisition and processing.

The experiment presents a systematic comparison of willingness to invest in assets with ambiguous return distributions. The design is based on the model in Dow and Werlang (1992) and deviates from standard approaches of measuring ambiguity attitudes with pairwise choices: Here we use revealed uncertainty premia as an indicator for ambiguity aversion. 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. To understand the impact of learning, we study investment decisions across two information conditions: one in which subjects make their decisions with specified probabilities; and another in which subjects must engage in belief updating before investing.

The structure of the experiment can therefore be summarized as a 2x2 design that allows for comparing decisions across two dimensions. A first dimension varies the degree of uncertainty by comparing decisions under risk versus ambiguity. The second and main dimension distinguishes between situations in which information about return distributions is released at once and those in which information is processed sequentially, requiring belief updating.

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. The control treatment "No Learning" (**NL**) investigates the relation between different types of uncertainty and investment decisions when belief updating is not required. This treatment verifies that our measures indeed reflect uncertainty premia. The treatment of main interest is "Learning" (**L**), which examines uncertainty premia when investors receive information gradually over time and engage in belief updating. More specifically, participants learn the distribution across two stages: they first receive information about a prior distribution and then observe an additional signal. The experiment is designed such that treatments NL and L should generate identical quotes under FBU, but substantially different quotes under MLU.

The data analyses of aggregate quotes favor FBU over MLU, but mask substantial heterogeneity in updating behavior. The majority of subjects (54%) chose quotes that conformed with the qualitative predictions of FBU. These subjects considered mainly the type of quote (ask vs. bid) when setting their quotes, exhibiting thereby significant ambiguity premia relative to risk. Another substantial group of subjects (34%) chose quotes in line with MLU: Depending on the signal they submitted extreme high or low quotes with minimum spreads. A remainder of subjects (11%) did not react to the signal information and exhibited similar ambiguity and risk premia.

The paper is organized as follows. Section 2 briefly discusses the closest literature. Section 3 presents the stylized decision model and the theoretical predictions. Section 4 describes the implementation of the experiment. The results are presented in Section 5, and Section 6 concludes.

2 Related Literature

The paper relates to a growing literature on belief updating under ambiguity. While some recent experimental work (e.g., Epstein and Halevy, 2019; Fryer et al., 2019; Kellner et al., 2019; Shishkin and Ortoleva, 2020; Liang, 2020) study learning with ambiguous information, we focus here on the updating of ambiguous priors with unambiguous signals.[3]

Another strand of the literature studies updating with a sequence of signals (Corgnet et al., 2013; Moreno and Rosokha, 2016; Baillon et al., 2018; De Filippis et al., 2018; Li and Wilde, 2019). De Filippis et al. (2018) study the updating of private signals when ambiguity about a prior belief emanates from social learning, and provide evidence for a generalization of MLU. Within the framework of recursive expected utility Moreno and Rosokha (2016) find deviations from Bayesian updating under ambiguity, while Li and Wilde (2019) do not for the majority of their subjects. Our setting separates from theirs in two ways: First, we consider the basic updating process with a single signal. Second, we focus on a specific ambiguity model–*maxmin expected utility*, for which FBU and MLU rules have been axiomatized. Corgnet et al. (2013) and Baillon et al. (2018) study learning in a financial market setting. Corgnet et al. (2013) finds no evidence of under- or overreaction or other ambiguity effects. Relatedly, Baillon et al. (2018) investigate whether belief updating interferes with ambiguity attitudes. Our approach is very different: We presume that ambiguity attitudes

[3]Liang (2020) also provides structural estimates on the updating of ambiguous priors that are in line with our general conclusion despite a different definition of uncertainty premia.

are generally stable, allowing us to make inferences about updating rules.

Other experimental studies investigate two properties of Bayesian updating–dynamic consistency and consequentialism. In their dynamic variants of the Ellsberg three-color experiment Dominiak et al. (2012) and Bleichrodt et al. (2018) find consequentialism to be less violated than dynamic consistency. Our experimental design assumes consequentialism and cannot speak to dynamic consistency as we elicit a single decision per subject. The dynamic Ellsberg game was originally presented in Cohen et al. (2000) who in their experiment also enforce consequentialism to differentiate between the FBU and MLU rules. They, too, find heterogeneity in updating behavior. In their study the preponderance of FBU over MLU approximates a ratio of 2:1. The current paper reproduces a very similar ratio in a different framework with prices rather than choices, highlighting the consequences of these two updating rules for trading activity.

3 A stylized decision problem

3.1 Optimal investment under EU and MEU

We first present the decision problem in the simple case where no learning occurs. In Section 3.2 we then discuss how introducing learning affects these predictions. For exposition we assume risk-neutrality throughout Section 3.

3.1.1 Expected Utility

Consider the payoff-relevant state space $S = \{V_L, V_H\}$, with $0 < V_L < V_H$. We denote $\pi = Pr(V = V_H)$ the probability for the high-value state, with $1 - \pi = Pr(V = V_L)$ as the complementary probability.

The agent is endowed with cash W_0 and has the opportunity to invest in *a single unit* of an asset with uncertain value $V \in S$. She 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 expected payoff of buying and selling are given by $(W_0 + E[V] - p)$ and $(W_0 + p - E[V])$, respectively.

The agent buys the asset if

$$E[V] - p \geq 0. \tag{1}$$

Analogously, she sells short (henceforth sell) the asset if

$$p - E[V] \geq 0. \tag{2}$$

It is thus optimal to set $a^* = b^* = E[V] = V_L + \pi(V_H - V_L)$. The risk-neutral agent buys at prices below her expected valuation, sells at prices above it. A risk-averse agent chooses a strictly positive spread between bid and ask, with $b^* < E[V]$ and $a^* > E[V]$ (the simple proof is in Appendix A.1).

3.1.2 Maxmin Expected Utility

We model ambiguity in π with multiple priors. The present argumentation therefore follows Dow and Werlang (1992), but with maxmin expected utility (MEU - Gilboa and Schmeidler, 1989) instead of Choquet expected utility.[4]

Let $\Pi = \{\pi | \pi_l \leq \pi \leq \pi_h\}$ denote the set of success probabilities. An MEU agent buys if

$$\min_{\pi \in \Pi} (E_\pi[V] - p) \geq 0. \tag{3}$$

She sells if

$$\min_{\pi \in \Pi} (p - E_\pi[V]) \geq 0. \tag{4}$$

An MEU agent considers the worst possible expected payoff, which implies that she considers different probability distributions for the two actions of buying and selling. Optimizing leads to two separate quotes $b^* = E_{\pi_l}[V] = V_L + \pi_l(V_H - V_L)$ and $a^* = E_{\pi_h}[V] = V_L + \pi_h(V_H - V_L)$, with $b^* < a^*$.

Figures 1 illustrates Dow and Werlang (1992)'s no-trade theorem under MEU in contrast to EU: The expected payoff functions of ambiguity-averse buying and selling strategies are shifted downwards, relative to the case of EU. Given 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. Hence, the ambiguity-averse seller sells at higher prices, while the ambiguity-averse buyer displays a lower willingness to pay. In between, there is

[4]FBU has been axiomatized for MEU but not for Choquet expected utility (Pires, 2002).

a range of prices at which buyer and seller do not agree on trade.[5,6]

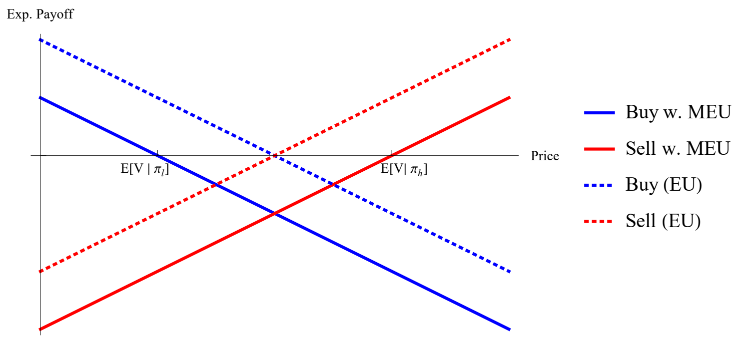

Note: Expected payoff of a buy and a sell as a function of the price for risk-neutral EU (dashed lines) and MEU (solid lines) agents.

Figure 1: Expected payoff of buy and sell with EU versus MEU

Under risk aversion, MEU predicts wider spreads than *subjective expected utility* (henceforth SEU, Savage 1954). Note that under SEU optimal bid and ask quotes for a given utility function are monotonically increasing in π. It follows that an MEU agent with a particular concave utility function and belief set Π will choose wider spreads than an SEU agent with the same utility function and any belief $\pi \in \Pi$ where $|\Pi| > 1$. In that case, the increase in spreads relative to SEU stems from the MEU agent considering multiple and more extreme probability measures.

3.2 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 correct with probability $q = P(s = \vartheta_L | V = V_L) = P(s = \vartheta_H | V = V_H)$. We denote the prior and posterior beliefs with $\mu = Pr(V = V_H)$ and $\rho = Pr(V = V_H | s, \mu)$, respectively. To formalize the difference between the two main information conditions, we will differentiate between the final probability π in the treatment without

[5]Under risk neutrality the general prediction of trade frictions holds for different models of ambiguity aversion (e.g., Choquet expected utility (CEU), maxmin expected utility (MEU), α-maxmin expected utility (α-MEU), smooth preferences), but the magnitude of spreads may differ across models.

[6]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).

learning, and the prior and posterior probabilities (μ, ρ) in the treatment with learning.

3.2.1 Bayesian updating

Applying Bayes' rule to a single prior belief μ yields a single posterior belief ρ. Analogously to equations (1) and (2), the agent buys if

$$E_\rho[V|s] - p \geq 0 \tag{5}$$

and sells if

$$p - E_\rho[V|s] \geq 0. \tag{6}$$

It is then optimal to quote the bid and ask $b^* = a^* = E_\rho[V|s] = V_L + \rho(V_H - V_L)$. That is, the agent adjusts the quotes to information but holds a zero spread before and after information. More generally, optimal quotes are the same with and without belief updating if and only if $\rho = \pi$. This holds under risk aversion as well, since the risk-averse EU agent will hold the same non-zero spread for the same belief value, regardless of final beliefs being exogenously given (π) or endogenously updated (ρ).

3.2.2 Full Bayesian updating

With an ambiguous prior optimal quotes depend on the way of updating beliefs. The literature has proposed various updating rules (Dempster, 1967; Jaffray, 1992; Epstein and Le Breton, 1993; Maccheroni et al., 2006; Epstein and Schneider, 2007; Hanany and Klibanoff, 2007; Hanany et al., 2009; Klibanoff et al., 2009; Siniscalchi, 2011), but so far only FBU (Fagin and Halpern 1991; Jaffray 1992, axiomatized by Pires 2002) and MLU (Dempster 1967; Shafer 1976, axiomatized by Gilboa and Schmeidler 1993; Ghirardato et al. 2008) have been axiomatized within the framework of MEU. These two main concepts make maximum opposite predictions with respect to the spread.

Let $\mathscr{M}$ be the set of priors and Π^B be the set of posteriors obtained by applying Bayes' rule to all priors $\mu \in \mathscr{M}$. Agents with multiple priors apply FBU when they update prior by prior to end up with a set of posteriors. Unless $q = 1$, Π^B is not a singleton and FBU does not fully eliminate ambiguity.

The optimization of an MEU agent is then analogous to (3) and (4). She buys if

$$\min_{\rho \in \Pi^B} (E_\rho[V|s] - p) \geq 0. \tag{7}$$

She sells if

$$\min_{\rho \in \Pi^B} (p - E_\rho[V|s]) \geq 0. \tag{8}$$

Thus, $b^* < a^*$ if $|\Pi^B| > 1$. The MEU agent chooses a non-zero spread both before and after the updating process. More importantly, it follows that FBU leads to the same optimal quotes and spreads than in the no learning framework if and only if $\Pi^B = \Pi$.

3.2.3 Maximum likelihood updating

Under MLU the set of posterior beliefs Π^B is obtained by applying Bayes' rule to a selected set of priors $\mu^* \in \mathscr{M}^* = \arg\max_{\mu \in \mathscr{M}} \ell(\mu|s)$, where $\ell(\mu|s)$ represents the likelihood of a prior μ given signal s.

Thus, 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, $|\mathscr{M}^*| = 1$, consequently $|\Pi^B| = 1$. For instance, after observing a high signal ϑ_H the agent assigns the highest likelihood to the highest prior $\mu_h = \sup \mathscr{M}$. Hence, the agent postulates a single posterior whenever a single prior maximizes the likelihood of having generated the informative event.

Equations (7) and (8) continue to hold, but because Π^B is here a singleton, signals eliminate any perception of ambiguity. The MLU agent will therefore adjust her belief to one of the two extremes, depending on the signal being high or low, and choose equal bid and ask $b^* = a^* = E_\rho[V|s]$ with $\rho = \rho(\mu^*, s) \in \{\rho_l(\mu_l, \vartheta_L), \rho_h(\mu_h, \vartheta_H)\}$.

In short, for a given set of priors $\mathscr{M}$ FBU and MLU generate different sets of posteriors Π^B, which in turn affects the choice of quotes and spreads. In our setting, the no-trade theorem of Dow and Werlang (1992) does no longer hold with MLU. Here, a fundamental difference between FBU and MLU is that different factors determine the ranking of states. With FBU the ranking of states depends on the long or short position (Mukerji and Tallon, 2001). An agent using MLU, on the other hand, ranks the states according to her information.

4 Experimental design

4.1 Treatment No Learning (NL)

In treatment No Learning (**NL**), subjects do not have to engage in belief updating. It therefore provides the benchmark for the average bid-ask spreads for ambiguous and risky assets.

Treatment **NL** consists of 20 independent rounds. At the beginning of each round, subjects receive information about π, learning in particular whether or not π is ambiguous. This information is visualized with an "urn A", which contains 100 balls in a mixture of red and blue balls. Subjects are told that to determine the asset value, the computer will draw a ball (henceforth "value ball") from urn A: the asset will take 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 1 for the chosen parameters) and is shown to the subjects. That is, subjects learn π for risky prospects by observing the exact number of red and blue balls in urn A. When π 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, that could potentially be red or blue, 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. Figure 2 shows examples of urn A with unambiguous and ambiguous distributions.

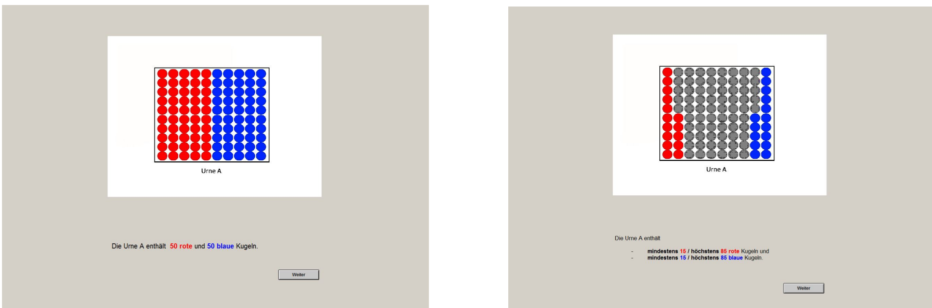

(a) Risky prospect **(b)** Ambiguous prospect

Figure 2: Examples for visualization of probability distribution with urn A.

The computer determines the asset value V by drawing a ball from urn A according to π in risky rounds; 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 start every round with an endowment of cash W_0 and, after seeing the information about π, quote both a bid and an ask ($b, a \in [V_L, V_H]$, $b < a$) on a second, separate screen. The computer then draws a random price p in $[V_L, V_H]$. The comparison between the subject's bid and ask quotes and the price determines whether a buy, a sell or no trade took place.

The round ends with subjects receiving feedback about the true asset value V, the random price p, the type of trade that took place (buy, sell, no trade) and their respective payoff in that specific round.

4.2 Treatment Learning (L)

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

In the first stage, subjects receive information about a 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. The signal corresponds to the color of another ball (henceforth "signal ball") that is drawn from a second urn. The composition 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.

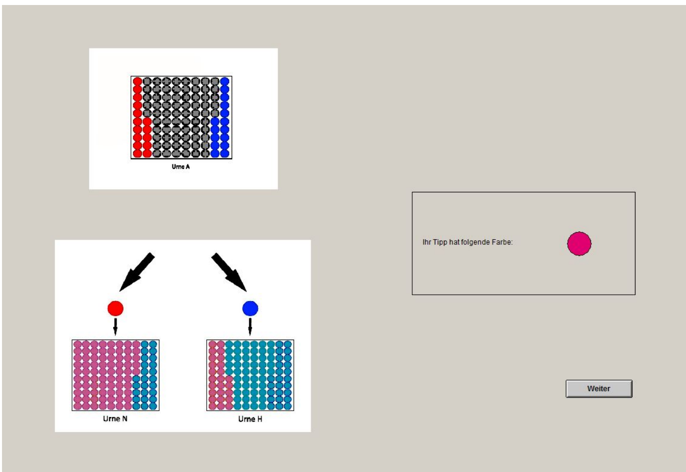

Figure 3: Example for an additional signal at the second stage.

Figure 3 depicts an example of the screen at the second stage. In the upper left corner, a figure with the composition in urn A reminds the subjects of the prior distribution; the graph below illustrates again the correlation between value and signal ball. The right side of the screen conveys the additional information by showing the color of the signal ball. 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).

Subjects then submit their bid and ask quotes on a separate screen, and at the end of a round receive feedback about the price, the asset value, the type of trade and their payoff.

4.3 Experimental procedures

The computerized experiment was run in the laboratory of Technical University Berlin and WZB Berlin Social Science Center.[7] In total, 110 and 112 students participated in treatments NL and L, respectively. Each treatment was run with five sessions of approximately 22 subjects.

The experiment consisted of two parts: the decision game and an additional part to elicit control measures of subjects' general attitudes toward risk, uncertainty and ambiguity (see Appendix Section C). The decision game started once all participants had read the instructions and had responded correctly to a comprehension test.

In the decision game subjects started each round with a cash endowment of $W_0 = 100$ experimental currency units (ECU). The asset could take either the value $V_L = 0$ or $V_H = 100$ ECU.

The set of possible probability values was chosen to be parsimonious in order to have more replications (i.e., more observations) for the comparison between treatments. In each treatment 14 out of 20 rounds were with unambiguous asset distributions, the remaining six were with ambiguous distributions. 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 1). 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 $\rho(s = \vartheta_H) \in [.35; .95]$).

The main feature of the experimental design was to set $\Pi^B = \Pi$ under ambiguity. That is, the objective sets of probabilities, $[.05; .65]$ and $[.35; .95]$, in NL were chosen to equal the set of posterior beliefs under FBU in L. This allows

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

Table 1: Chosen values for the probability π and the prior μ with corresponding Bayesian posterior ρ

	No Learning	Learning		
			$\rho(s = \vartheta_L)$	$\rho(s = \vartheta_H)$
Risk	$\pi = .05$	$\mu = .05$	$\rho = .02$	$\rho = .14$
	$\pi = .15$	$\mu = .15$	$\rho = .05$	$\rho = .35$
	$\pi = .35$	$\mu = .35$	$\rho = .15$	$\rho = .62$
	$\pi = .50$	$\mu = .50$	$\rho = .25$	$\rho = .75$
	$\pi = .65$	$\mu = .65$	$\rho = .38$	$\rho = .85$
	$\pi = .85$	$\mu = .85$	$\rho = .65$	$\rho = .95$
	$\pi = .95$	$\mu = .95$	$\rho = .86$	$\rho = .98$
	$T_R = 7 \times 2 = 14$	$T_{RI} = 7 \times 2 = 14$		
	Probability	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 $[.15, .85]$.

for comparing bids and asks for the same, final dispersion in probabilities, when information on the distribution was provided at once versus across two stages.

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 two out of the five 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 already set before making any decisions, but they learned which round was chosen only after finishing the decision game. The instructions as well as the computer screen emphasized accordingly that hedging across rounds made no sense once the payoff-relevant round was set.

Earnings consisted of a show-up fee (5 EUR), plus two thirds of earnings in the randomly drawn round in the decision game plus one third of earnings in a randomly chosen task for the elicitation of attitudes. 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, 20.50 EUR for approximately 100 minutes.[8]

4.4 Hypothesis

Hypothesis 1 *With FBU and* $|\Pi| > 1$*,*

$$\begin{aligned} &E[a\ |\Pi^B] = E[a\ |\Pi] \\ &E[b\ |\Pi^B] = E[b\ |\Pi] \\ &\texttt{with}\ E[a-b|\Pi^B] = E[a-b|\Pi] > 0 \\ &\texttt{and}\ \Pi^B = \Pi. \end{aligned}$$

Remember, the experiment is designed such that $\Pi^B = \Pi$, i.e., the set of

[8] Six sessions were run in 2016, the remaining 4 in 2020. At the end of these last four sessions, subjects received an additional "surprise" lump-sum payment of 2 EUR to match the increased expected payment of the laboratory.

FBU beliefs in treatment L match the set of specified probabilities in treatment NL. If subjects' beliefs conform predominantly with FBU, receiving information gradually will have no major impact in that final beliefs–and therefore quotes–will, on average, be similar across treatments NL and L. Subjects with MLU beliefs, on the other hand, will consider a substantially different set of beliefs $\Pi^B \subset \Pi$, choosing smaller spreads and more extreme quotes.

Two main premises lead to our hypothesis. The first and most important one is $E[a-b|\Pi] > 0$. That is, subjects choose an ambiguity premium consistent with MEU preferences. Allowing for risk aversion we will posit

$$E[a-b|\Pi] > E[a-b|\pi] \geq 0, \quad \forall\, \pi \in \Pi \texttt{ and } |\Pi| > 1. \tag{9}$$

In other words, subjects choose wider spreads under ambiguity than under risk. Note that (9) identifies ambiguity aversion that cannot be rationalized with SEU. More specifically, bid-ask pairs for a set Π that are more divergent than bid-ask pairs chosen at any $\pi \in \Pi$–i.e., at all unambiguous probability values in the same set–are interpreted as evidence consistent with ambiguity aversion but not SEU. Our second premise is that subjects are Bayesian. We discuss the validity of these premises in Section 5.1 and Appendix Section B.2.

5 Results

We present the results as follows: First, we confirm the premise that the spread reflects an uncertainty premium. Second, we investigate the main hypothesis that subjects' quotes are consistent with FBU beliefs.

5.1 Spread as an Uncertainty Premium

We use the control treatment NL to expose the empirical relationship between spreads and uncertainty at three levels.

First, spreads that capture a risk premium should increase in the variance of returns, i.e., when the success probability π converges to 50%. As can be seen in Figure 4a, this is confirmed by the median spread that matches the risk of investing: it is hump-shaped in the success probability, with a maximum at a probability of 50%. Overall, subjects chose a median spread of 5 ECU in risky rounds (see Appendix Table B1).

Second, MEU implies wider spreads under ambiguity relative to risk. This conjecture is also compatible with the data: median spreads for prospects with ambiguous probabilities are two to three times as high as with unambiguous

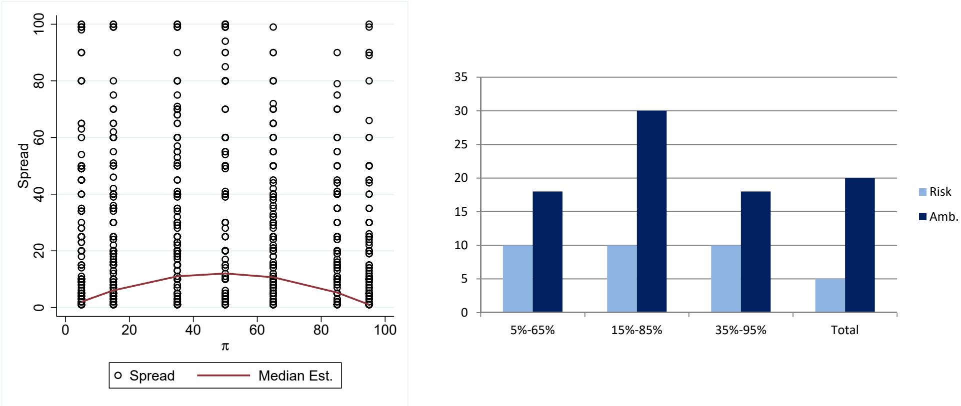

(a) Risk **(b)** Risk vs. Ambiguity

Figure 4: Median spreads as a function of π

probabilities, reflecting non-negligible ambiguity aversion (see Figure 4b). Note that this size in spreads cannot be rationalized with SEU since subjects chose substantially wider spreads under ambiguity than under risk for all $\pi \in [\pi_l, \pi_h]$.

Third, more ambiguity generates wider spreads. Figure 4b also shows that subjects chose wider spreads with a wider range for ambiguous probabilities: For $\pi \in [15\%, 85\%]$ the median spread amounted to 30 ECU, which was 12 ECU more ($p = 0.004$ in median test) than in the other ambiguous rounds with $\pi \in [5\%, 65\%]$ or $\pi \in [35\%, 95\%]$ (see Appendix Table B1 for more details).

In a nutshell, subjects reacted to uncertainty by choosing wider spreads, corroborating our premise that spreads reflect an uncertainty premium.

5.2 Testing FBU

Given the prevalent ambiguity premium in treatment NL, we next investigate the main hypothesis that subjects are full Bayesian updaters. Remember, the experiment was designed such that under FBU subjects would choose the same spreads and the same level of quotes in the ambiguous rounds of treatments NL and L for the same set of final beliefs (i.e., $\Pi = \Pi^B$).

Descriptive Statistics. The descriptive statistics are inconclusive. While the distribution of spreads is consistent with FBU, the distribution of quotes lends some support to MLU theory.

Subjects chose the same average spread for the same degree of uncertainty. As can be seen in Figures 5a and 5b, the distribution of spreads is not statistically different ($p = 0.94$ in the Kolmogorov-Smirnov test), despite a slightly larger median spread in treatment L (18 in NL vs. 20 in L, $p = 0.89$ in median test). This first result is at odds with MLU theory.

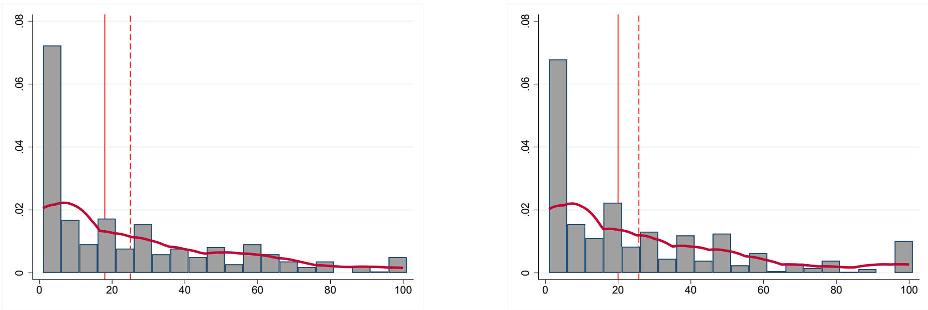

(a) Ambiguous rounds in NL **(b)** Ambiguous rounds in L

Note: Median and mean spreads are represented by the vertical solid and dashed line, respectively. The solid graph corresponds to the estimated kernel density.

Figure 5: Spreads for ambiguous prospects where $\Pi = \Pi^B$

To examine the distribution of quotes we summarize each subject's bid and ask quotes by computing the midquote $\frac{a+b}{2}$, which is the midpoint between chosen bid and ask quotes. Appendix Figures (B1a) to (B1d) show the distribution of midquotes separately for low and high probability values. It can be readily seen that midquotes substantially differed between the two treatments ($p = 0.002$ in Kolmogorov-Smirnov test). We are particularly interested in knowing whether quotes were more extreme in treatment L, which would rather be consistent with MLU. To this end, we measure how extreme quotes were with ($|\frac{a+b}{2} - 50|$), the absolute deviation of midquotes to the center 50. On average, midquotes deviated more from the center in treatment L (by 20 units compared to 15 units in treatment NL, $p < 0.001$).

FBU versus MLU at the Aggregate Level. Given the inconclusive descriptive statistics we resort next to a structural estimation to investigate which of the two updating rules fits the data better. We derive the test assuming MEU under ambiguity.

Let A and S be dummy variables that take the value one if the quote is an ask and if the signal is high, respectively. Define $\rho_l(\mu)$ and $\rho_h(\mu)$ as the Bayesian update of the prior μ after a low and a high signal, respectively. For exposition we normalize quotes and beliefs to the same range [0,100]. With this simplified notation MEU predicts $a, b \in \{\inf \Pi^B, \sup \Pi^B\}$, and thus allows us to focus on the two extreme beliefs that are relevant. We can then rewrite FBU quotes as the Bayesian update of the two extreme priors μ_l and μ_h.

$$\begin{aligned} q^{FBU} &= \rho_l(\mu_l) + [\rho_l(\mu_h) - \rho_l(\mu_l)] \cdot A + [\rho_h(\mu_l) - \rho_l(\mu_l)] \cdot S \\ &+ [(\rho_h(\mu_h) - \rho_h(\mu_l)) - (\rho_l(\mu_h) - \rho_l(\mu_l))] \cdot A \cdot S \end{aligned} \tag{10}$$

As can be easily seen from rearranging (10), the trading position A determines whether the FBU quote is a function of μ_l or μ_h. In other words, whether Bayesian updates of a high or low prior flows into decision-making is governed by whether the quote is an ask or a bid. The prediction in (10) contrasts to the quotes based on MLU beliefs:

$$q^{MLU} = \rho_l(\mu_l) + [\rho_h(\mu_h) - \rho_l(\mu_l)] \cdot S \tag{11}$$

For MLU quotes the prior under consideration is solely determined by the signal.

Plugging in the parameter values specified in Table 1 leads to $q^{FBU} = 5 + 60 \cdot A + 30 \cdot S$ and $q^{MLU} = 5 + 90 \cdot S$. Thus, in our design FBU quotes are mainly determined by A, the trading position, while S is the sole predictor under MLU. We therefore estimate the following linear hierarchical model, using the quotes for ambiguous assets in treatment L:

$$\begin{aligned} &\boldsymbol{q}_i = \boldsymbol{X}_i\boldsymbol{\beta} + \boldsymbol{Z}_i\boldsymbol{u}_i + \boldsymbol{\varepsilon}_i \\ &\texttt{with } \boldsymbol{u_i} \sim N(\boldsymbol{0}, \boldsymbol{\Gamma}),\ \boldsymbol{\varepsilon_i} \sim N(\boldsymbol{0}, \boldsymbol{\Sigma}) \end{aligned} \tag{12}$$

The dependent variable $\boldsymbol{q_i}$ is the $(J \times 1)$ vector of all $J = 12$ quotes that subject i submitted for ambiguous prospects. The $(J \times k)$ matrix of predictors $\boldsymbol{X}_i$ differs by model, with elements $\boldsymbol{x}'_{ij} = \begin{pmatrix} 1 & S_{ij} & A_{ij} \end{pmatrix}$ and $\boldsymbol{x}'_{ij} = \begin{pmatrix} 1 & S_{ij} \end{pmatrix}$ under FBU and MLU, respectively. The $(J \times k)$ design matrix $\boldsymbol{Z}_i$ allows for subject-specific random effects $\boldsymbol{u}_i$ of dimension $(k \times 1)$. The $(J \times 1)$ vector $\boldsymbol{\varepsilon}_i$ contains the decision-specific errors. The covariance matrices $\boldsymbol{\Gamma}$ and $\boldsymbol{\Sigma}$ may take any form, allowing for a correlation between subject i's various random effects, and between subject i's decision errors.[9]

Table 2 summarizes the main results of the estimation. To discriminate between these two theories, we test whether the chosen quotes in the data are more governed by the trading position (ask vs. bid) or the signal (high vs. low).[10]

[9]For instance, by design bid and ask quotes of the same rounds are not independent given $b < a$.

[10]Five subjects who did not have any variation in the signal were dropped from the estimation.

Table 2: RESULTS OF THE HIERARCHICAL MODEL

Model	FBU	MLU
Signal	39.01	38.33
	(2.49)	(2.56)
Ask	25.97	
	(2.08)	
Cons.	14.43	27.75
	(1.14)	(1.47)
N	107	107
Total ($N \times J$)	1284	1284
Log Pseudo-Lik.	-5599.41	-5975.71

Note: Robust standard errors in parentheses.

Both the signal and the trading position are important determinants of chosen quotes. Hence, the BIC approximation to the Bayes factor yields $1.61e+157$, showing strong support in favor of the FBU model.[11] Note, however, that subjects put significantly more weight on the signal than on the trading position ($p = 0.0006$). The qualitative prediction of FBU according to which the trading position has more weight than the signal is therefore not satisfied.

FBU versus MLU at the Individual Level. The analysis at the aggregate level masks the substantial heterogeneity in the data, which led to the inconclusive descriptive statistics.[12] This heterogeneity becomes apparent in Appendix Figures B1c and B1d that display more variation in quotes in treatment L compared to treatment NL. A finite mixture model helps uncover homogeneous types of decision making in this heterogeneous mass. The mixture model classifies subjects into three types, according to their approach to setting quotes:

[11] We use Wagenmakers' (2007) formula $BF_{01} \approx exp(\Delta BIC_{10}/2)$.

[12] This heterogeneity does not seem to be driven by a lack of probabilistic sophistication. Section B.2 in the Appendix shows that Bayesian inference cannot be rejected in risky rounds.

$$\begin{aligned}
&\texttt{Type 1 (FBU)}: && q_{ij} = \beta_0 + \beta_1 S_{ij} + \beta_2 A_{ij} + \varepsilon_{1it} \\
&\texttt{Type 2 (MLU)}: && q_{ij} = \alpha_0 + \alpha_1 S_{ij} + \varepsilon_{2it} \\
&\texttt{Type 3 (LI)}: && q_{ij} = \gamma_0 + \gamma_1 A_{ij} + \varepsilon_{3it} \\
& && i = 1, ..., n\ , j = 1, ...12 \\
& && \varepsilon_{1it} \sim N(0, \sigma_1^2), \quad \varepsilon_{2it} \sim N(0, \sigma_2^2), \quad \varepsilon_{3it} \sim N(0, \sigma_3^2),
\end{aligned}$$

The first type corresponds to subjects who consider both the signal and the trading position when submitting a quote. We call these subjects, whose quoting strategy is in line with FBU, FBU subjects. The second type of subjects consider the signal only– as predicted with MLU. With a third type we also include the possibility that subjects are insensitive to changes in the likelihood of V_H, but have an uncertainty premium. In that case, their quotes do not react to the information in the signal, but only to the trading position. We label these subjects as likelihood insensitive (LI).

The log pseudo-likelihood function is given by [13]:

$$\mathscr{L} = \sum_{i=1}^{n} lnL_i \tag{13}$$

with

$$L_i = p_1 \prod_{j=1}^{12} \phi(q_i|\texttt{FBU}) + p_2 \prod_{j=1}^{12} \phi(q_i|\texttt{MLU}) + (1 - p_1 - p_2) \prod_{j=1}^{12} \phi(q_i|\texttt{LI}) \tag{14}$$

where p_1 and p_2 represent the proportions of types 1 and 2 in the population, respectively. The function $\phi(\cdot)$ corresponds to the normal density function.

[13] The true likelihood may depart from our specification as there may be residual correlation between bid and ask quotes of the same round. We use this pseudo-likelihood function to identify types of subjects. We then check the robustness of the results and inferences by re-estimating the model within types and accounting for residual correlation.

Table 3: Results of the finite mixture model*

	FBU	MLU	LI
Signal	26.95	65.49	
	(2.08)	(2.65)	
Ask	38.72		15.92
	(2.73)		(3.53)
Cons.	16.56	13.46	29.24
	(1.53)	(1.42)	(1.51)
σ_ε	21.19	15.05	15.16
N	58	38	11

* *Note*: This table shows the results of a random effects GLS regression with cluster-robust standard error at the subject level after classifying subjects into types. The results of the finite mixture model with the pseudo-likelihood function are almost identical and can be found in Appendix Table B4.

Table 3 shows the results of the finite mixture analysis. The majority of subjects (54%) acted like FBU subjects. They took into account both the signal and the trading position, while putting more emphasis on the latter ($p = 0.002$ in one-sided test). These subjects chose higher spreads under ambiguity than under risk (median spread of 35 in ambiguous rounds versus 15 in risky rounds, $p < 0.001$ in median regression).

A second substantial group of subjects (35%) behaved in accordance with MLU: They primarily focused on the signal information when choosing their quotes. On average, they chose a minimal spread of 1, both in ambiguous and risky rounds ($p = 1$ in median regression). A small remainder of subjects (10%) neglected the signal information. This group was generally more willing to sell than to buy, and chose, on average, a spread of 11.5 ECU, which was similar to the spread chosen in risky rounds ($p = 0.66$ in median regression). Appendix Section B.3 provides more details to the distributions of quotes and spreads within each group.

We note that the estimates within each group conform with the qualitative, but deviate from the exact quantitative predictions of their respective model (i.e., $q^{FBU} = 5 + 60 \cdot A + 30 \cdot S$ and $q^{MLU} = 5 + 90 \cdot S$). Surely, the quantitative predictions under MEU are extreme. In between these extremes, a generalization of MLU with $1 < |\mathscr{M}^*| < |\mathscr{M}|$ describes the case where subjects update a nonsingleton subset of priors that they consider to be likely. It is conceivable that the less extreme estimates in the FBU group conform with a generalization of MLU. Yet, in this setting generalizations of MLU are more difficult to separate from FBU with some errors in Bayesian updating. This experiment specifically targets the qualitative contrast between FBU and MLU within the framework of MEU, and its results are interpreted as such.

6 Conclusion

This experiment reveals that updating under ambiguity is heterogeneous. A majority of subjects' decisions were in line with FBU, but there was also a non-negligible share of subjects whose decisions conformed with MLU. These findings suggest that ambiguity effects cannot generally be detached from the information condition. The same degree of ambiguity may lead to different decisions, depending on whether or not information is released gradually and on who is processing this information.

These findings emphasize the role of learning for generating differential beliefs in ambiguous settings. This is in line with Bossaerts et al. (2010), who argue that heterogeneity has important implications for markets, which are, therefore, not best described by a representative agent. Heterogeneity in updating behavior, though, may impact markets differently than heterogeneity in ambiguity attitudes. In particular, gradual information flow may reinforce disparities in decisions under ambiguity. For instance, asset pricing might be determined by extreme updaters if those who update cautiously refrain from trading.

Of course, this experiment compares two extreme paradigms within the framework of MEU. This stylized structure allows us to identify different pricing strategies and their respective decision factors. Whether an alternative updating rule–possibly applied to different preference models– captures the heterogeneity in decision factors better than the two paradigms remains an open question for future research.

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, Tim Christensen, Peter Cramton, Erik Eyster, Guillaume Frechette, Dan Friedman, Itzhak Gilboa, Yoram Halevy, Philippe Jehiel, Peter Klibanoff, Matthias Lang, Andrew Schotter, as well as two referees, the Editor and the participants of seminars at several universities. All errors are mine. I thank the Humboldt-Universität zu Berlin, the DIW Berlin and the Center for Experimental Social Science at New York University 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. *American Economic Review 101*(2), 695–723.

Baillon, A., H. Bleichrodt, U. Keskin, O. L'haridon, and C. Li (2018). The effect of learning on ambiguity attitudes. *Management Science 64*(5), 2181–2198.

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

Bleichrodt, H., J. Eichberger, S. Grant, D. Kelsey, and C. Li (2018). A Test of Dynamic Consistency and Consequentialism in the Presence of Ambiguity.

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.

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.

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 (2018). Non-Bayesian updating in a social learning experiment.

Dempster, A. P. (1967). Upper and lower probabilities induced by a multivalued mapping. *Annals of Mathematical Statistics 38*(2), 325–339.

Dominiak, A., P. Duersch, and J. P. Lefort (2012). A dynamic Ellsberg urn experiment. *Games and Economic Behavior 75*(2), 625–638.

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.

Epstein, L. G. and Y. Halevy (2019). Hard-to-Interpret Signals.

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.

Fagin, R. and J. Y. Halpern (1991). A new approach to updating beliefs. *Uncertainty in Artificial Intelligence 6*, 347–374.

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

Fryer, R. G., P. Harms, and M. O. Jackson (2019). Updating Beliefs when Evidence is Open to Interpretation: Implications for Bias and Polarization. *Journal of the European Economic Association 17*(5), 1470–1501.

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

Gilboa, I. and M. Marinacci (2016). Ambiguity and the Bayesian Paradigm. In H. Arlo-Costa, V. Hendricks, and J. van Benthem (Eds.), *Readings in Formal Epistemology*, Chapter 21, pp. 385–439. Springer Graduate Texts in Philosophy.

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 (2019). Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study.

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.

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

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.

Jaffray, J.-Y. (1992). Bayesian updating and belief functions. *IEEE Transactions on Systems, Man, and Cybernetics 22*(5), 1144–1152.

Kellner, C., M. T. Le Quement, and G. Riener (2019). Reacting to ambiguous messages: An experimental analysis.

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

Li, W. and C. Wilde (2019). Belief Formation and Belief Updating Under Ambiguity: Evidence from Experiments. *SAFE Working Paper* (251).

Liang, Y. (2020). Learning from unknown information sources.

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

Machina, M. J. and M. Siniscalchi (2014). Ambiguity and ambiguity aversion. *Handbook of the Economics of Risk and Uncertainty 1*, 729–807.

Marinacci, M. (2015). Model Uncertainty. *Journal of the European Economic Association 13*(6), 1022–1100.

Moreno, O. M. and Y. Rosokha (2016). Learning under compound risk vs. learning under ambiguity– an experiment. *Journal of Risk and Uncertainty 53*(2-3), 137–162.

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.

Savage, L. J. (1954). *The Foundations of Statistics* (John Wiley & Sons, revised and enlarged by Dover Publications ed.). New York.

Shafer, G. (1976). *A Mathematical Theory of Evidence* (Princeton University Press ed.). Princeton.

Shishkin, D. and P. Ortoleva (2020). Ambiguous Information and Dilation: An Experiment.

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.

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review 14*(5), 779–804.

A Mathematical appendix

A.1 Bid-ask spread due to 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 Results

B.1 Reactions to ambiguity

B.1.1 Descriptive statistics

Table B1: Median and mean spread for various ranges of probabilities.

$\pi \in$	$[5\% - 65\%]$	$[15\% - 85\%]$	$[35\% - 95\%]$	Total obs.	
		Median		Mean	Median
Risk	10	10	10	18.11(1.39)	5
Amb.	18	30	18	28.51(1.95)	20
Diff.	-8***	-20***	-8**	-10.40***	-15***
N	1320	1320	1320	2200	

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.

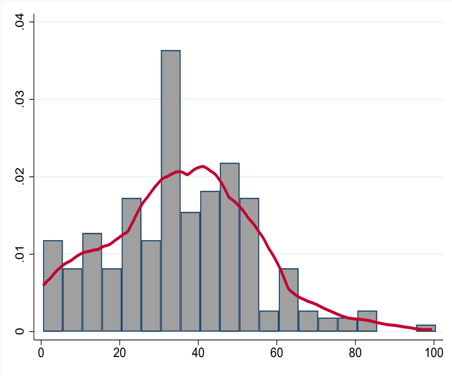

(a) $\pi \in [.05, .65]$

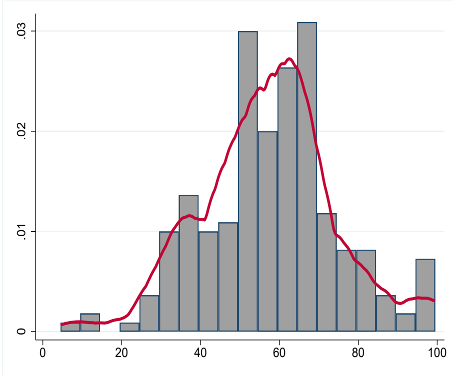

(b) $\pi \in [.35, .95]$

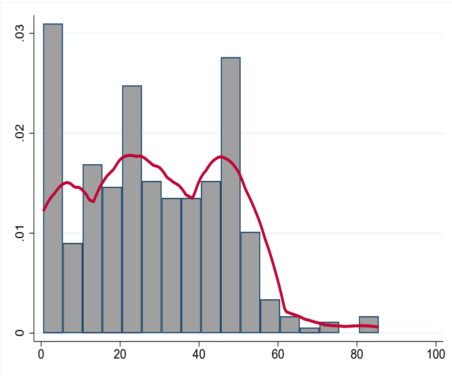

(c) $\rho \in [.05, .65]$

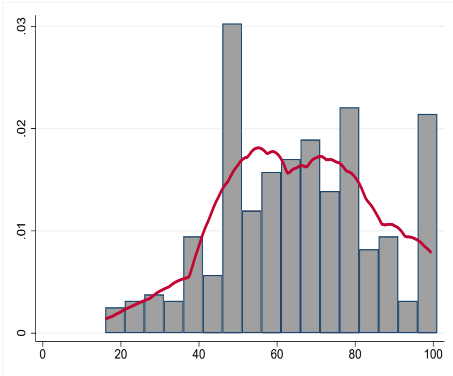

(d) $\rho \in [.35, .95]$

Note: 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.

Figure B1: Distribution of midquotes

Figures B1a and B1b show the distribution of midquotes in treatment NL for a low (left) and a high (right) range of probabilities. Distributions are fairly symmetric and centered around 35 and 58 ECU, respectively. In treatment L, quotes are more heterogeneous with more mass approximately around (0,20,50) after a low signal (Figure B1c) and around (50,80,100) after a high signal (Figure B1d). Since the conditional probability for a correct signal is $q = .75$, the mass points close to 25 and 75 suggest base-rate neglect as a possible explanation. However, as discussed in Section B.2., we find no evidence of base-rate neglect in decisions with unambiguous return distributions, in which subjects adjusted their quotes to the risky prior.

B.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):

$$\begin{cases} b_i = (1 - RP_b) \cdot E[V|\tilde{\tau}] + \epsilon_{i,b} \\ a_i = (1 + RP_s) \cdot E[V|\tilde{\tau}] + \epsilon_{i,s}, \end{cases} \tag{15}$$

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. Subject-specific decision errors in buying and selling are captured with $(\epsilon_{i,b}, \epsilon_{i,s})$. 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 $\tau = \rho$.[14] The coefficient α regulates the curvature of the function. The parameter β determines the inflection point of the curve.

The probability weighting function is, in general, inverse s-shaped, reflecting a general over-weighting of small and under-weighting of high probabilities.

[14] 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.

Table B2: PROBABILITY WEIGHTING FUNCTION AND RISK PREMIA

	NL		L	
β	0.7765	(.0313)	0.7947	(.0418)
α	0.7706	(.0535)	0.6906	(.046)
RP_s	0.0114	(.0239)	0.0205	(.0235)
RP_b	0.2479	(.0210)	.2403	(.0262)

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

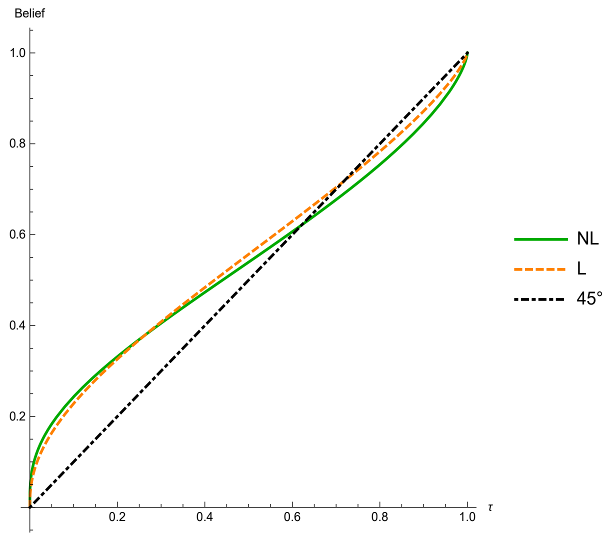

Figure B2: Estimated probability weighting function for unambiguous probabilities in NL & L.

The functions do not differ between the two treatments. That is, subjects reacted to unambiguous marginal probabilities in the same way as to unambiguous Bayesian posteriors. Assuming a stable relation between decisions and probabilities, Bayesian inference cannot be rejected.

We also investigate the prevalence of base-rate neglect in treatment L. To this end we estimate the relation between midquotes for risky prospects and the information contained in signals and risky priors. Table B3 shows the estimates of Grether (1980)'s regression model $ln(\psi) = c + \gamma_1\ ln(\frac{\mu}{1-\mu}) + \gamma_2\ ln(LR(s)) + u$, where ψ denotes the odds ratio of midquotes and $LR(s)$ refers to the likelihood ratio of the signal. In treatment L the weights assigned to signal and prior information do not significantly differ from each other, showing no evidence of base-rate neglect. Figure B3 visualizes how in treatment L midquotes for risky prospects increase with unambiguous priors. We also estimate the analog model for treatment NL with $ln(\psi) = c + \gamma_1\ ln(\frac{\pi}{1-\pi}) + u$, and find that the weight assigned to the specified probability values (π in NL or μ in L) is the same across treatments.

Table B3: Bayesian Inference

	NL		L	
$\hat{\gamma}_1$	0.8035	(.047)	0.7795	(.056)
$\hat{\gamma}_2$	-		0.9057	(.091)
c	-0.0573	(.054)	-0.1057	(.058)
R^2	0.5185		0.5064	
N	1540		1568	

Note: CRSE in parentheses. Estimates of γ are not significantly different from each other.

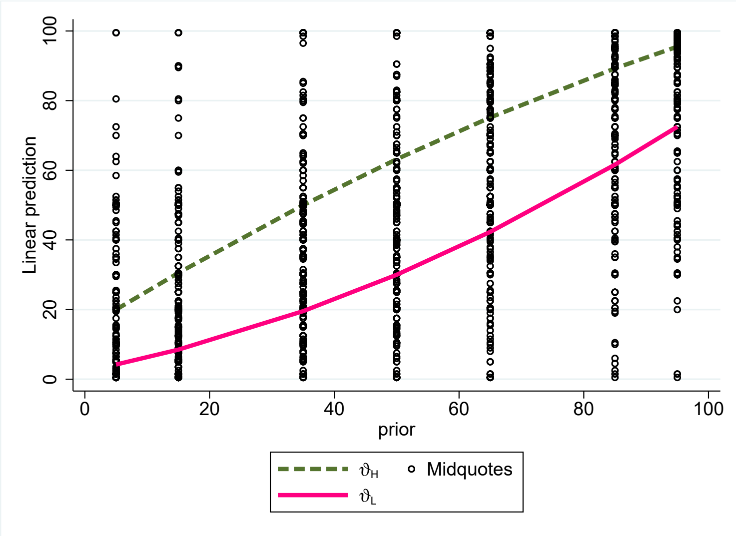

Note: The dashed and solid lines correspond to median estimates after subjects receive a high and a low signal, respectively.

Figure B3: Midquotes as function of risky priors.

B.3 Results of the finite mixture model

Table B4 shows the direct estimates of the finite mixture analysis.

Table B4: Estimates of the finite mixture model

	FBU	MLU	LI
Signal	26.36	65.64	
	(2.02)	(2.17)	
Ask	38.69		16.90
	(2.20)		(3.37)
Cons.	17.00	13.31	29.20
	(1.46)	(1.20)	(1.99)
σ_ε	21.99	16.19	15.55

Figures B4a to B4c show the bid-ask pairs chosen within each group of the finite mixture model. Deviations from the 45-degree line translate into higher spreads. Quotes from MLU subjects are visibly different from the ones of FBU subjects by being more extreme and exhibiting small spreads, whereas quotes from LI subjects are similar after low and high signals.

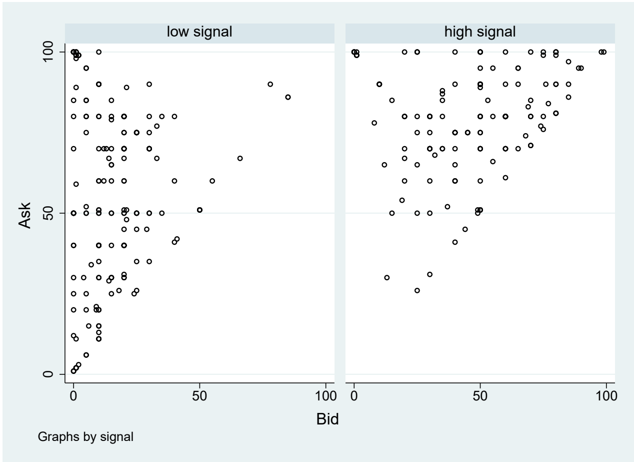

(a) FBU

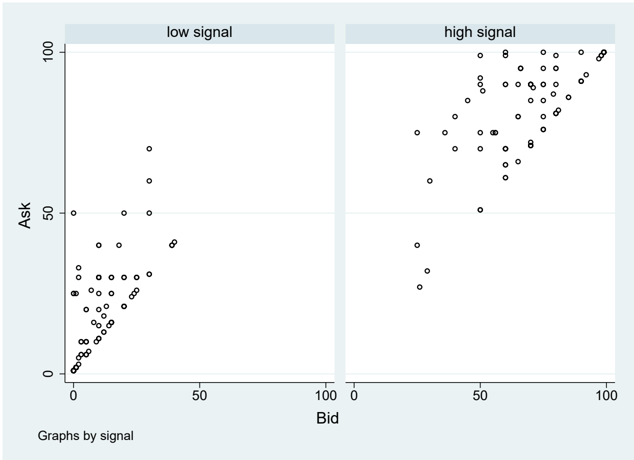

(b) MLU

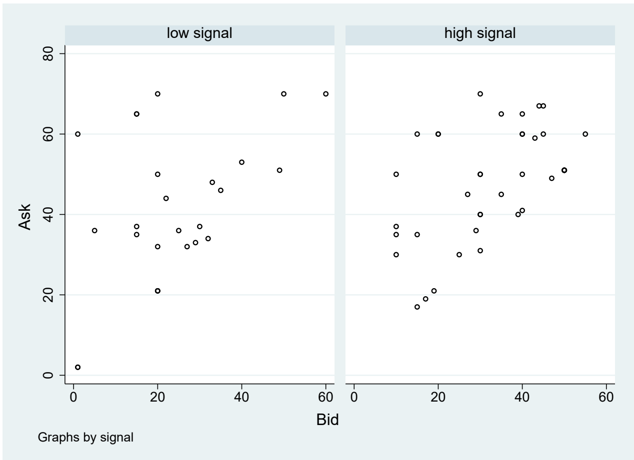

(c) LI

Figure B4: Bid-ask pairs after high and low signals

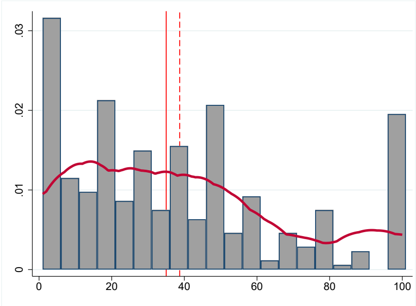

(a) FBU

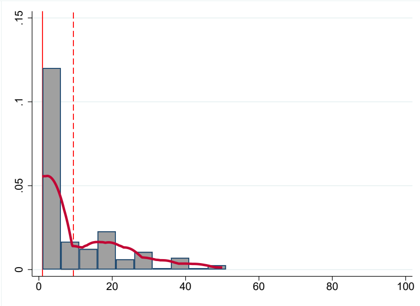

(b) MLU

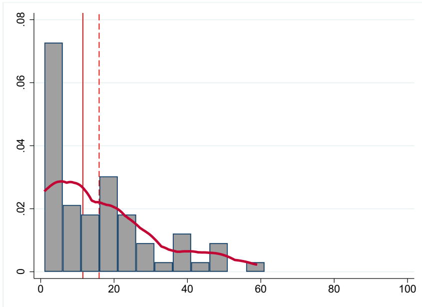

(c) LI

Figure B5: Distribution of spreads by groups

The same conclusions are obtained by looking at the distribution of spreads (see Figures B5a to B5c). FBU subject chose generally wider spreads than MLU subjects.

C Elicitation of attitudes toward uncertainty

We elicited control measures of attitudes toward different types of uncertainty 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, we 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. We find that attitudes toward risk and uncertainty correlate with decisions in the main part of the experiment, but attitudes toward ambiguity do not. We find the majority of subjects to be ambiguity seeking. This casts some doubts on the reliability of the elicitation procedure. For instance, the lack of randomization may bias attitudes in a systematic direction, 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.

C.1 Risk attitudes

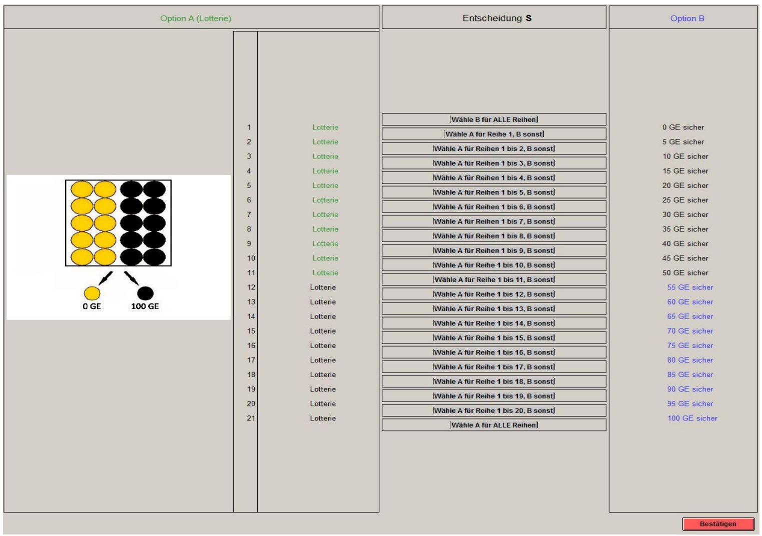

Figure C1: Example of computer interface in Part 2

Risk preferences were elicited with a multiple price list task akin to Abdellaoui et al. (2011) and Gillen et al. (2019). 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 C1 depicts the computer interface for the first measurement with lottery $(100, 0.5; 0)$.

C.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.

C.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.

C.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.

C.3 Results

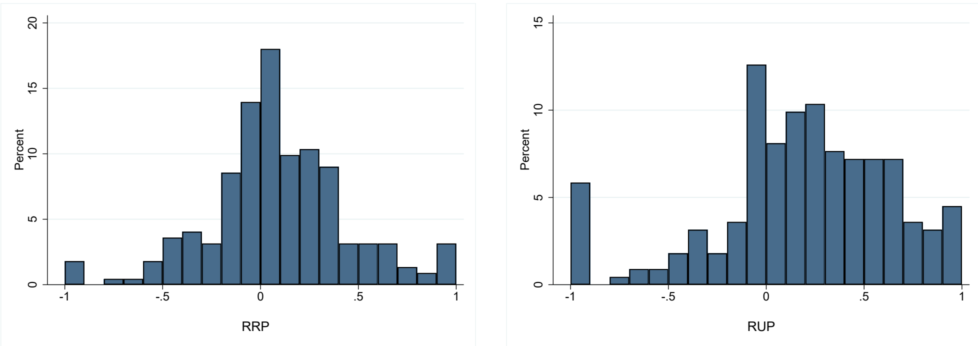

(a) RRP **(b)** RUP

Figure C2: Relative risk and uncertainty premia

We 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 (C2a) and (C2b) 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, 61.71% and 68.42% of all subjects chose a positive risk and uncertainty premium, respectively. A share of 22.07% had neither a positive risk nor uncertainty premium, 16.22% 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 average positive uncertainty premia. Only 22.97% of subjects were ambiguity averse but a surprisingly large fraction of 68.47% 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 (see Table C1). 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 -0.03).

Table C1: Average spreads for different categories of elicited attitudes

		mean spread	med. spread
RRP	< 0	15.12	5
		(0.64)	
	> 0	20.28***	10
		(0.58)	
RUP	< 0	20.48	10
		(1.17)	
	> 0	30.15***	20
		(0.96)	
δ	< 1	23.99	15
		(1.46)	
	$= 1$	28.90	25
		(3.20)	
	> 1	27.91	20
		(0.96)	

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