Centre de Recherche en Économie et Management
Center for Research in Economics and Management

University of Caen

University of Rennes 1

Working Paper

Learning under ambiguity: An experiment using initial public offerings on a stock market

Aurélien Baillon,
Han Bleichrodt,
Umut Keskin

Erasmus University Rotterdam , Erasmus School of Economics, the Netherlands

Olivier L'Haridon

CREM UMR CNRS 6211 and GREGHEC, University of Rennes 1, France

Chen Li

Erasmus University Rotterdam , Erasmus School of Economics, the Netherlands

August 2013 - WP 2013-31

Learning under ambiguity:
An experiment using initial public offerings on a stock market

Aurélien Baillon[a], Han Bleichrodt[a], Umut Keskin[a], Olivier L'Haridon[b], Chen Li[a]

[a]Erasmus School of Economics, Rotterdam, the Netherlands.

[b]University of Rennes 1-Crem, Greg-HEC, Rennes, France.

13 August 2013

Abstract

This paper studies the effect of learning new information on decision under uncertainty. Using ambiguity models, we show the effect of learning on beliefs and ambiguity attitudes. We develop a new method to correct beliefs for ambiguity attitudes and decompose ambiguity attitudes into pessimism (capturing ambiguity aversion) and likelihood insensitivity. We apply our method in an experiment using initial public offerings (IPOs) on the New York Stock Exchange. IPOs provide a natural decision context in which no prior information on returns is available. We found that likelihood insensitivity decreased with information, but pessimism was unaffected. Subjects moved in the direction of expected utility with more information, but significant deviations remained. Subjective probabilities, corrected for ambiguity attitudes, were well calibrated and close to market data.

Keywords: ambiguity, learning, updating, neo-additive weighting.

ACKNOWLEDGEMENTS**:** Ferdinand Vieider and Peter P. Wakker gave helpful comments. Han Bleichrodt's research was made possible by a grant from the Netherlands Organization for Scientific Research (NWO).

1. Introduction

This paper studies experimentally how decision makers change their behavior in the face of new information about uncertain events. If objective probabilities are unknown, the traditional approach in economics is to assume that the decision maker can assign subjective probabilities to events and behaves according to expected utility. In expected utility, subjective probabilities are used as decision weights and are updated using Bayes' rule.

While people change their beliefs upon the arrival of new information and these updated beliefs have predictive value (Hamermesh 1985, Smith et al. 2001), they systematically deviate from Bayesianism.[1] This has economic implications. For example, a recent study by Ju and Miao (2012) showed that Bayesian learning is at odds with various dynamic asset-pricing phenomena.

In a fundamental contribution, Ellsberg (1961) challenged the very existence of subjective probabilities. Ellsberg's paradox undermined not only the validity of subjective expected utility (which had already been done by Allais's paradox), but also the more general notion of probabilistic sophistication (Machina and Schmeidler 1992). In reaction to Ellsberg's paradox, new ambiguity models of decision under uncertainty have been developed (for overviews see Wakker 2010, Gilboa and Marinacci forthcoming). Most of these ambiguity models allow for the possibility that decision weights differ from subjective probabilities. The decision weights reflect not only people's beliefs but also the confidence they have in these beliefs and their aversion towards ambiguity. The ambiguity models capture an intuition expressed by Keynes (1921):

"The magnitude of the probability of an argument…depends upon a balance between what may be termed the favourable and the unfavourable evidence; a new piece of evidence

[1] See Grether (1980), El-Gamal and Grether (1995), Charness and Levin (2005), Hoffman et al. (2011), Poinas et al. (2012). Psychologists and behavioral economists have found many updating biases, including under- and overconfidence (Griffin and Tversky 1992), conservatism (Phillips and Edwards 1966), representativeness (Kahneman and Tversky 1972), availability (Tversky and Kahneman 1973), and confirmatory bias (Rabin and Schrag 1999) and suggested heuristic decision models to explain these biases.

which leaves this balance unchanged also leaves the probability of the argument unchanged. But it seems that there may be another respect in which some kind of quantitative comparison between arguments is possible. This comparison turns upon a balance, not between the favourable and the unfavourable evidence, but between the *absolute* amounts of relevant knowledge and relevant ignorance respectively" [p.71].

In Keynes' words, learning of new evidence changes both the balance of evidence (people`s beliefs) and the total amount of evidence (the amount of ambiguity). Under expected utility, the amount of ambiguity plays no role and learning only affects beliefs. In the ambiguity models new information changes both beliefs and ambiguity attitudes and they make it possible to better understand the effects of learning on behavior This raises the question of how decision weights are updated.[2] While several papers have approached this question from a theoretical angle and different updating rules have been proposed,[3] there is a dearth of empirical evidence on how decision weights are actually updated.[4] This motivated our paper.

We study the updating of decision weights in the context of a general preference model (Miyamoto 1988, Luce 1991, Ghirardato and Marinacci 2001) that contains most ambiguity models as special cases. We then present a simple method to disentangle beliefs and ambiguity attitudes. Our method describes a decision maker's ambiguity attitude by two indices, one reflecting his pessimism (capturing ambiguity aversion) and the other his sensitivity to changes in likelihood. Insensitivity to likelihood has been interpreted as a cognitive bias, and is probably most affected by new information.

[2] In the literature the expression "updating of non-Bayesian beliefs" is sometimes used. To emphasize that beliefs may differ from subjective probabilities under non-expected utility we use the term updating of decision weights.
[3] See Gilboa and Schmeidler (1993), Epstein (2006), Eichberger et al. (2007), Epstein and Schneider (2007), Hanany and Klibanoff (2007), Eichberger et al. (2010), Eichberger et al. (2012).
[4] Cohen et al. (2000) and Dominiak et al. (2012) experimentally studied updating under ambiguity but consider situations in which decision makers receive information that an event cannot occur. In our study decision makers accumulate evidence in favor or against events.

The separation of beliefs and ambiguity attitudes makes it possible to study whether people behave more in line with expected utility when they receive more information. If new information decreases pessimism and likelihood insensitivity then the decision maker will move towards expected utility maximization. This would be compatible with a common view in economics that learning and more information decrease irrationalities caused by deviations from expected utility (Myagkov and Plott 1997, List 2004, van de Kuilen and Wakker 2006).

Expected utility is still widely seen as the normative standard for decision under uncertainty. However, it is also well known that people deviate from expected utility. The discrepancy between the normative and descriptive status of expected utility makes it desirable to adjust preference measurements for deviations from expected utility. There is a large literature on correcting utility measurements for deviations from expected utility (McCord and de Neufville 1986, Wakker and Deneffe 1996, Delquié 1997, Bleichrodt et al. 2001). Our paper complements this literature by showing how the measurement of beliefs can be corrected for deviations from expected utility.

We applied our method in an experiment, where subjects traded options with payoffs contingent on the performance of (anonymous) initial public offerings (IPOs). IPOs make it possible to study the effect of new information in a natural decision context (rather than in a more contrived context using urns) for which no prior information is available. We found that the arrival of new information reduced subjects' likelihood insensitivity, but not their pessimism. Beliefs were well-calibrated after correction for ambiguity attitude and reflected aggregate market behavior. Subjects` behavior moved in the direction of expected utility when they obtained more information, but significant deviations remained.

2. Theoretical framework

Decision model

A decision maker faces uncertainty about the outcome he will receive at time T. This uncertainty is modeled through a finite *state space* S_T where the subscript T denotes that the uncertainty will be resolved at time point T. The state space contains all possible *states of the world* s. Only one state occurs, but the decision maker does not know which one. *Events* are subsets of S_T. The decision maker chooses between *binary acts*, denoted by $x_E y$, giving money amount x if event E occurs at time T and money amount $y \le x$ otherwise.

The decision maker's information about previous resolutions of uncertainty up to time $t < T$ is formalized by his *history set* $h_t = (s_1, \ldots, s_t)$, where $s_j \in S_j$ for all $1 \le j \le t$, and S_j denotes the state space representing the uncertainty at time j. Complete absence of information is denoted h_0.We assume that $S_t = S_T = S$ *for all* $t = 1, \ldots, T$. In other words, the same states are available at different points in time. The decision maker's beliefs may vary over time as new information arrives. The decision maker's preferences are represented through a *history-dependent preference relation* $\succcurlyeq_t$ where the subscript t indicates that preferences depend on the history h_t . The relations $\succ_t$ and $\sim_t$ are defined as usual. A real-valued function V_t *represents* $\succcurlyeq_t$ if for all binary acts $x_E y, v_F w$, $x_E y \succcurlyeq_t v_F w$ iff $V_t(x_E y) \ge V_t(v_F w)$.

The Bayesian approach assumes that preferences $\succcurlyeq_t$ are represented by *expected* utility: $x_E y \mapsto P_t(E)U(x) + \big(1 - P_t(E)\big)U(y)$, with U a utility function defined over outcomes and P_t the subjective probability measure given h_t . In expected utility, new information, which expands the history set from h_t to h_v , with $v > t$, affects beliefs (subjective probabilities) but leaves utility unchanged. Updating takes place in the belief part of the representation and "tastes" (utility) are not influenced by new information regarding past events. The assumption of constant utility is also common in the theoretical literature on the updating of decision weights under non-expected utility (e.g. Epstein 2006, Eichberger et al. 2007, Epstein and Schneider 2007). We will also adopt it in this paper.

To account for deviations from expected utility, we will assume a *binary rank-dependent utility* (RDU) model (Miyamoto 1988, Luce 1991Ghirardato and Marinacci 2001), which includes most ambiguity models as special cases. Examples include contraction expected utility (Gajdos et al. 2008), maxmin expected utility (Gilboa and Schmeidler 1989), alpha-maxmin expected utility (Ghirardato et al. 2004), Choquet expected utility (Schmeidler 1989), and prospect theory (Tversky and Kahneman 1992). Under binary RDU $\succcurlyeq_t$ can be represented by

$$x_E y \mapsto W_t(E)U(x) + \big(1 - W_t(E)\big)U(y), \qquad (1)$$

with *U a real-valued function* unique up to level and unit and W_t a unique *weighing function,*[5] which may be non-additive but satisfies $W_t(\varnothing) = 0$, $W_t(S_T) = 1$ and $W_t(A) \leq W_t(B)$ if $A \subseteq B$. The subscript t in W_t expresses that weights depend on the history just like P_t in the Bayesian approach.

Chateauneuf et al. (2007) introduced a tractable way to analyze decision weights W_t, *neo-additive weighting*, in which decision weights are a linear function on (0,1). For parameters a_t and b_t such that $a_t \leq 1$ and $a_t - 2 \leq b_t \leq 2 - a_t$, and for a probability measure P_t, neo-additive decision weights are defined as

$$\begin{aligned}
W_t(E) &= \tfrac{a_t - b_t}{2} + (1 - a_t)P_t(E) && \text{if } 0 < \tfrac{a_t - b_t}{2} + (1 - a_t)P_t(E) < 1,\\
W_t(E) &= 0 && \text{if } \tfrac{a_t - b_t}{2} + (1 - a_t)P_t(E) \leq 0, \text{ and} \qquad (2)\\
W_t(E) &= 1 && \text{if } \tfrac{a_t - b_t}{2} + (1 - a_t)P_t(E) \geq 1.
\end{aligned}$$

[5] Sometimes the term *capacity* is used instead of weighting function.

Neo-additive decision weighting assumes that the decision maker is *probabilistically sophisticated for a given history*: his decisions are consistent with a probability measure P_t. Because a_t and b_t differ across histories, the decision maker does not satisfy probabilistic sophistication in general and can deviate from it when comparing acts involving different histories.

Equation (1) with neo-additive weighting can be written as

$$x_E y \mapsto (1-a_t)\big[P_t(E)U(x) + \big(1-P_t(E)\big)U(y)\big] + \frac{a_t-b_t}{2}U(x) + \frac{a_t+b_t}{2}U(y). \quad (3)$$

Equation (3) is a linear combination of expected utility, the maximum utility $U(x)$, and the minimum utility $U(y)$. This expression helps to understand the intuition underlying the parameters a_t and b_t as we show next. We will refer to Eq. (3) as the *neo-additive model*.

Chateauneuf et al. 2007) imposed the stronger constraints $0 \le a_t \le 1$ and $-a_t \le b_t \le a_t$. We will call this the *natural case* because it ensures that decision makers are likelihood insensitive and assign positive weights to extreme outcomes ($-a_t \le b_t \le a_t$). Our constraints also permit likelihood oversensitivity and zero weights for extreme outcomes. This made it possible to include more subjects in our analyses.

Likelihood insensitivity

The parameter a_t in Eq. (3) reflects the weight that the decision maker gives to expected utility in his evaluation of acts. If a_t is equal to 0 then the decision maker behaves according to expected utility. The larger is a_t, the less weight the decision maker gives to expected utility and the more he concentrates on the maximum and minimum utility. In other words, the larger a_t the more the decision maker ignores the relative likelihoods of x and y. This can also be seen from Eq. (2), where the larger a_t the lower the weight given to $P_t(E)$.

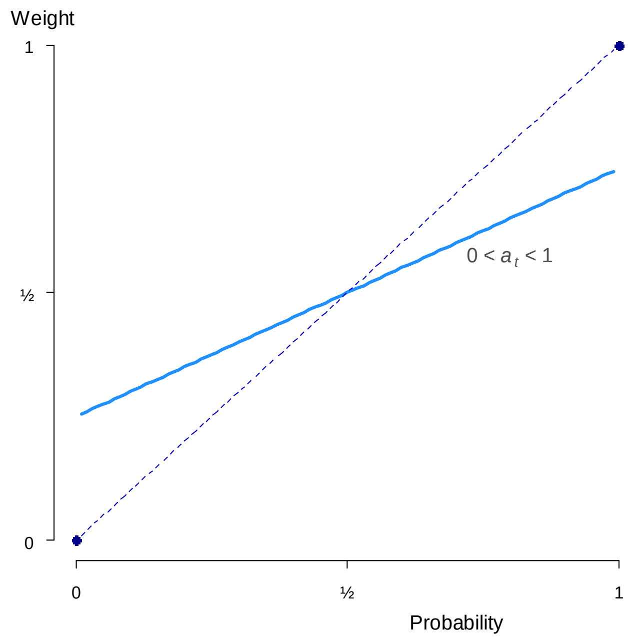

Figure 1. Likelihood insensitivity.
The figure shows the neo-additive weighting function with $a_t > 0$ and $b_t = 0$. The decision maker is insufficiently sensitive to changes in likelihood. The diagonal shows the weighting function when expected utility holds.

Figure 1 shows the effect of changes in a_t when b_t is held constant to 0. When $a_t = 0$, the decision maker behaves according to expected utility (dashed line). When a_t becomes more positive the slope of the decision weighting function becomes flatter and the decision maker is less sensitive to intermediate changes in likelihood. He does not perfectly discriminate between likelihood levels and differences between (non-extreme) decision weights are less than the differences between their underlying probabilities. This is called *likelihood insensitivity*. We take a_t as a *likelihood insensitivity index* with higher values of a_t corresponding with more likelihood insensitivity.

Empirical studies usually found more likelihood insensitivity for uncertainty than for risk (e.g. Kahneman and Tversky 1979, Kahn and Sarin 1988, Kilka and Weber 2001, Abdellaoui et al. 2005, Wakker 2010). There is also evidence that likelihood insensitivity is stronger for less familiar sources of uncertainty (Kilka and Weber 2001, Abdellaoui et al. 2011). We therefore expected that likelihood insensitivity is negatively related to the size of the history set (the amount of information).

Pessimism

Figure 2 shows that for a given value of a_t, increases in b_t shift the weighting functions downwards. Because the decision weights reflect the weight given to the best outcome, increases in b_t imply that the decision maker pays more attention to the worse outcome. We will interpret b_t as an index of *pessimism* with higher values indicating more pessimism, and negative values reflecting optimism An expected utility maximizer has $b_t = 0$. In the natural case, an extremely pessimistic decision maker, who only considers the worst outcome regardless of its likelihood, has $b_t = 1$. and an extremely optimistic decision maker, who only considers the best outcome has $b_t = -1$.

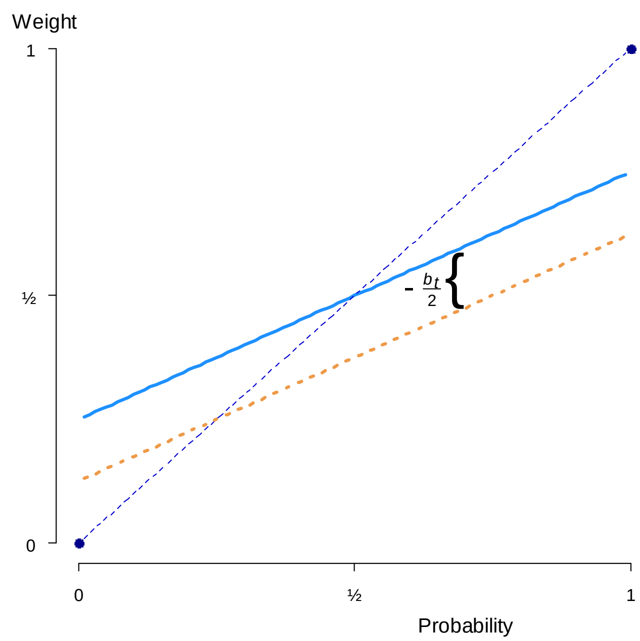

Figure 2. Pessimism.
The blue line corresponds to $a_t > 0$ and $b_t = 0$. The parallel dashed line keep a_t constant and increases b_t. The figure shows that increases in b_t shift the neo-additive weighting function downwards leading to an increase in pessimism.

Several studies found that pessimism decreased when the decision maker knew more about a source of uncertainty (Heath and Tversky 1991, Kilka and Weber 2001, Fox and Weber 2002, Di Mauro 2008, and Abdellaoui et al. 2011). These results suggest that pessimism will decrease when the history set becomes richer.

The effect of new information on beliefs on the one hand, and on likelihood sensitivity and pessimism on the other hand, illustrates that modern ambiguity theories capture Keynes' (1921) intuition about the weight and the balance of evidence. If new information changes the balance of evidence in favor of an event, the decision maker will update his beliefs accordingly. But this new information also changes the balance between the "absolute amounts of relevant evidence and relevant ignorance." Our approach captures this by allowing the decision maker to also update his weighting of subjective probabilities. The new information might induce the decision maker to rely more on his beliefs and become more sensitive to likelihood, with a_t tending to 0. In the next Section we will present a method to disentangle beliefs, pessimism, and likelihood insensitivity and to obtain beliefs that are corrected for ambiguity attitudes.

Multiple-prior interpretation of the neo-additive model

The above analysis is close to Choquet expected utility (Gilboa 1987, Schmeidler 1989) and prospect theory (Tversky and Kahneman 1992) where ambiguity attitudes are modeled through the decision weighting function. The multiple priors models take a different approach and model ambiguity through a set of priors C_t about the true probability measure P_t. Chateauneuf et al. (2007) showed that the neo-additive model also has a multiple-prior interpretation in the natural case. It can be rewritten as:

$$x_E y \mapsto alpha_t \, min_{\pi \in C_t}\left[\pi(E)U(x) + \left(1-\pi(E)\right)U(y)\right] +$$

$$(1-alpha_t) max_{\pi \in C_t}\left[\pi(E)U(x) + \left(1-\pi(E)\right)U(y)\right] \tag{4}$$

with $alpha_t = \frac{a_t + b_t}{2a_t}$ and $C_t = \{\pi: \pi(E) \geq (1-a_t)P_t(E), \text{ for all events } E\}$.

The set of priors C_t reflects the decision maker's perceived ambiguity; the larger the set of priors, the more ambiguity he perceives. The parameter $alpha_t$ reflects the decision maker's pessimism. Equation (4) is a linear combination of the lowest and the highest expected utility that the decision maker may obtain. The higher $alpha_t$, the more weight he gives to the lowest expected utility.

The set of priors is a function of a_t and P_t. If $a_t = 0$, the decision maker only considers P_t and there is no ambiguity. For positive a_t, he will also consider other probability measures. Thus a_t also measures the amount of ambiguity. This agrees with our previous interpretation of a_t as the cognitive component of decision under ambiguity. Increases in b_t, our measure of pessimism, lead to increases in $alpha_t$. However, $alpha_t$ also depends on a_t and, therefore, it is a different measure of pessimism than b_t.

Equation (4) is mathematically equivalent to Eq. (3) when a_t is positive. We cannot distinguish these interpretations and the reader can choose the interpretation that he likes best. However, the multiple-prior interpretation only holds in the natural case. Because several subjects had a negative a_t, we will use only Eq.(3) in the individual analyses. Because the mean value of a_t was positive we will analyze the aggregate data under both interpretations.

3. Measuring beliefs and ambiguity attitudes

We now explain how we identified a_t and b_t for different histories h_t. For each history h_t, we partitioned the state space S into three events. The events were defined by change in the price of stocks on a specific trading day. They were *Up*, the price goes up by at least 0.5%; *Middle*, the price varies by less than 0.5%; and *Down*, the price decreases by at least 0.5%. We also considered the event *MiddleUp* = *Middle* ∪ *Up*. For given payoffs $x > y$, we then elicited four certainty equivalents, $CE_{Up} \sim x_{Up}y$, $CE_{Middle} \sim x_{Middle}y$, $CE_{Down} \sim x_{Down}y$, and $CE_{MiddleUp} \sim x_{MiddleUp}y$. With the normalization $U(x) = 1$ and $U(y) = 0$, Eq. (1) implies that $U(CE_{Up}) = W_t(Up)$, $U(CE_{Middle}) = W_t(Middle)$, $U(CE_{Down}) = W_t(Down)$, and $U(CE_{MiddleUp}) = W_t(MiddleUp)$. The decision weights of an expected utility maximizer are equal to his subjective probabilities and, consequently, his subjective probabilities are equal to the utilities of his certainty equivalents. Thus under expected utility, $U(CE_{MiddleUp}) + U(CE_{Down}) = P_t(MiddleUp) + P_t(Down) = 1$. The utilities of the complementary events *MiddleUp* and *Down* should sum to 1. We will refer to this as *complementarity*. The neo-additive model allows for violations of complementarity:

$$U(CE_{MiddleUp}) + U(CE_{Down})$$
$$= \frac{a_t - b_t}{2} + (1 - a_t)P_t(MiddleUp) + \frac{a_t - b_t}{2} + (1 - a_t)P_t(Down) = 1 - b_t. \quad (5)$$

Equation (4) shows that the more pessimistic the decision maker, the lower the sum of $U(CE_{MiddleUp})$ and $U(CE_{Down})$. Hence, studying deviations of $U(CE_{MiddleUp}) + U(CE_{Down})$ from 1 allow us to identify the decision maker's degree of pessimism.

Under expected utility, the decision maker should also satisfy binary additivity: $U(CE_{Up}) + U(CE_{Middle}) - U(CE_{MiddleUp}) = P_t(Up) + P_t(Middle) - P_t(MiddleUp) = 0$. Under the neo-additive model, we obtain,

$$U(CE_{Up}) + U(CE_{Middle}) - U(CE_{MiddleUp}) = \frac{a_t - b_t}{2}. \quad (6)$$

As we know b_t from the test of complementarity, a_t can be uniquely identified.

The neo-additive model makes it possible to measure likelihood insensitivity and pessimism for any events if we can partition the state space S into three events and if we can measure utility. To measure utility we used the method of Abdellaoui et al. 2008), which we will explain below. Once we know a_t, b_t, and utility, we can also determine P_t. If a_t or b_t is unequal to zero, expected utility does not hold and measured beliefs will be non-additive. Our method takes this non-additivity into account and measures beliefs P_t that are corrected for ambiguity attitude.

4. Experiment

Subjects

The experiment was run at Erasmus University with 64 subjects (22 female). Subjects were either third year undergraduate students majoring in finance or graduate students in finance. Their average age was 24, ranging from 21 to 33. We used finance students because the experiment involved options. Finance students might better understand the experimental tasks and be more motivated to answer the questions. Each subject received a €5 show-up fee and, in addition, played out one of his choices for real using a procedure described below.

Procedure

The experiment was computer-run in small group sessions involving at most 3 subjects. Subjects first received instructions and then answered several questions to check their understanding of the experimental tasks. They could only proceed to the actual experiment after answering these questions correctly. The experimental instructions including the questions to check for subjects' understanding are in Appendix B.

As source of uncertainty we used the variation in the stock returns of IPOs (Initial Public Offerings) traded at the New York Stock Exchange (NYSE). IPOs are stocks that have just entered the market. We chose IPOs for two reasons. First, stock returns are a natural source of uncertainty unlike, for example, Ellsberg urns. Second, because IPOs are new on the market, there is no previous history of prices available and learning can occur.

We used data on 328 IPOs listed on the NYSE between 1 September 2009 and 25 February 2011. At the start of the experiment, each subject drew four numbers, which determined the stocks he would trade in. The identity of these stocks was revealed only after subjects had completed the experiment. Then we also explained subjects how they could verify the stock data on the internet.

Payoffs were determined by the performance of the stocks on the 21st trading day after their introduction on the NYSE. We defined four events: Up = $(0.5, \rightarrow)$, i.e. the stock goes up by more than 0.5% on the 21st trading day, Middle = $[-0.5, 0.5]$, the stock varies by at most 0.5% on the 21st trading day, Down = $(\leftarrow, -0.5)$, the stock goes down by more than 0.5% on the 21st trading day, and MiddleUp = $[-0.5, \rightarrow)$, the stock goes up by at least –0.5% on the 21st trading day. In what follows, we will refer to an option that pays x if event Up obtains as an *Up -option*. *Middle-, Down-*, and *MiddleUp options* are defined similarly. We used the variation in the stock instead of the absolute prices of the stocks to make sure that subjects had no information about the stocks and to avoid biases. Stocks with higher prices might attract more attention leading to biases in the elicited ambiguity attitudes.

There were three *informational conditions*, each involving a different history set. In the *no information* condition (history set h_0), subjects had no information about the stock returns. In the *one week* condition (history set h_5), subjects knew the daily returns of the stock during the first 5 trading days after its introduction. Finally, in the *one month* condition (history set h_{20}), subjects knew the stock returns during the first 20 trading days after its introduction.

Table 1: The 20 choice questions

Stock	***Condition***	***y***	***x***	***Option type***		***Stock***	***Condition***	***y***	***x***	***Option type***
1	No info	0	10	Up		3	1 week	0	20	Up
1	No info	10	20	Up		3	1 week	0	20	Middle
1	No info	5	20	Up		3	1 week	0	20	Down
1	No info	10	15	Up		3	1 week	0	20	MiddleUp
1	No info	0	5	Up		3	1 week	0	20	Middle
1	No info	0	20	Up		4	1 month	0	20	Up
2	No info	0	20	Up		4	1 month	0	20	Middle
2	No info	0	20	Middle		4	1 month	0	20	Down
2	No info	0	20	Down		4	1 month	0	20	MiddleUp
2	No info	0	20	MiddleUp		4	1 month	0	20	Down

The columns labeled "Stock" refer to the four different stocks. The questions for stock 1 were used to measure utility. The columns labeled "Condition" refer to the amount of information subjects received about the performance of the stock. Options were of the type $x_E y$ where the subject received €x if event E occurred and €y otherwise The columns "Option types" indicate event E.

We used choice lists to elicited the ask prices of 20 options, summarized in Table 1. Figure 3 gives an example of a choice list for a Middle-Up option. Subjects were told that they owned the option $x_E y$ and they were asked for each price on the choice list whether they wanted to sell the option. The choice lists consisted of 20 prices ranging from € $(y + z)$ to €x in increments of $z = € \frac{x-y}{20}$. The options corresponded with the stocks subjects had drawn at the start of the experiment. We used choice lists because previous research suggests that choice-based procedures lead to fewer inconsistencies than directly asking subjects for their certainty equivalents (Bostic et al. 1990, Noussair et al. 2004).

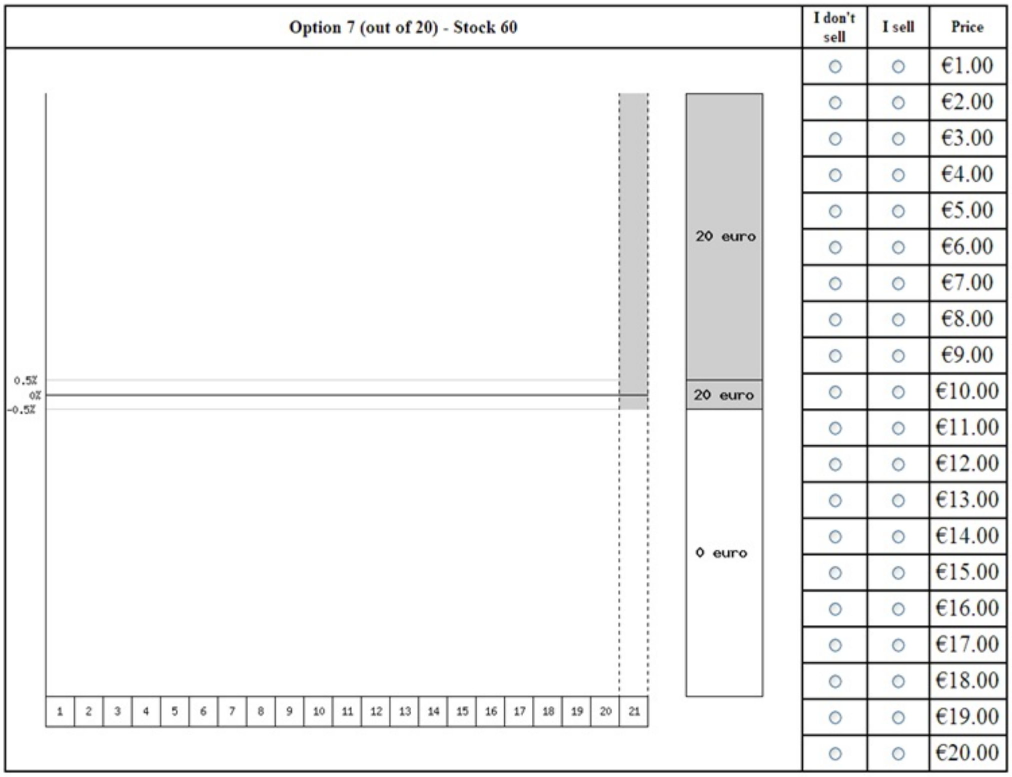

Figure 3. The choice lists used in the experiment. In this example the option pays €20 if event Middle-Up occurs on the 21st trading day after the introduction of the stock and €0 otherwise.

The 20 choices were divided into four groups (see Table 1). Group 1 consisted of six choices to measure utility. The questions in groups 2, 3, and 4 measured the effect of new information on ambiguity attitudes. For groups 3 and 4, we repeated one measurement to test the reliability of our measurements.

The utility questions (group 1) always came first. The order of the other groups was randomized to avoid that the effect of new information was confounded by a better understanding of the task. We had to use different stocks in each group. If we had used options on the same underlying stock, then subjects who had, for instance, received information on the performance of the stock in the first month would have used this information in the no information and in the one week conditions. We also randomized the order in which subjects faced the different options within each group.

Incentives

We used a random incentive system. At the end of the experiment, subjects threw a twenty-sided die twice. The first throw selected the option and the corresponding choice list and the second throw selected the line of that choice list that was played out for real. In the selected line, we implemented the choice that the subject had made during the experiment. If the subject had chosen to sell, we paid him the price. If he had chosen not to sell, we played out the option $x_E y$ and he received €*x* if event *E* had occurred on the 21st trading day and €*y* otherwise.

Analysis

All analyses used the original measurements. The repeated measurements were only used to test for consistency. For a given history, we excluded subjects whose the certainty equivalent of $20_{MiddleUp}0$ was less than the maximum of the certainty equivalents of $20_{Up}0$ and $20_{Middle}0$ minus €1. These subjects violated monotonicity. We used the €1 margin because preferences are typically imprecise and €1 was about equal to the median absolute deviation in the consistency tests. Three subjects violated monotonicity in the no information condition, eight in the one week condition, and nine in the one month condition. In the paired comparisons between conditions, we excluded subjects who violated monotonicity in at least one of the two conditions. To test for robustness we also analyzed the data excluding all subjects who violated monotonicity at least once. The results were similar.

Utility was measured using the method of Abdellaoui et al. (2008). We selected history h_0 and elicited certainty equivalents CE_j for the six binary acts $x_{j_{Up}}y_j, j = 1, \dots, 6$, the first entries of Table 1. By binary RDU:

$$U(CE_j) = W_0(Up)U(x_j) + (1 - W_0(Up))U(y_j).^{6} \tag{7}$$

We assumed a power utility function, i.e., $U(x) = x^{\beta}$ if $\beta > 0$, $U(x) = \ln(x)$ if $\beta = 0$, and $U(x) = -x^{\beta}$ if $\beta < 0$. The power family is widely-used in economics and decision theory and generally fits the data well (Stott 2006).

We used nonlinear least squares to estimate $W_0(Up)$ and β in (7). We then substituted β in Eqs. (5) and (6) to derive a_t and b_t and the subjective probabilities. Dividing all money amounts by the maximum payoff €20 scales the power utility function such that $U(20) = 1$ and $U(0) = 0$.

To account for the stochastic nature of subjects' preferences, we also estimated β, a_t, b_t and P_t using structural maximum likelihood estimation. Let $\theta = (a_0, b_0, \Delta a_{1week}, \Delta b_{1week}, \Delta a_{1month}, \Delta b_{1month}, P_t(Up), P_t(Middle), \sigma, \varepsilon)$ denote the vector of estimated model parameters for $t = 0$, 1 week, and 1 month. For t=1 week and t=1 month, $\Delta a_t = a_t - a_0$ and $\Delta b_t = b_t - b_0$. The parameter σ denotes error and ε denotes a tremble. Following Hey and Orme (1994), we assumed a Fechner error specification, which is widely used in stochastic choice under risk (e.g. Bruhin et al. 2010, Conte et al. 2011). Let $d_{ijtr} = 1(-1)$ denote the binary indicator that subject i chose to keep (sell) the option j with history t for binary choice r. The likelihood contribution for subject i facing option j with history t in choice r is:

$$l(d_{ijtr}|\theta, \sigma, \varepsilon) = (1 - \varepsilon)\Phi\left(d_{ijtr} \frac{W_t(E_{jt})U(x_{jt}) + (1 - W_t(E_{jt}))U(y_{jt}) - U(y_{jt} + (r-1)z_{jt})}{\sigma} \right) + \varepsilon/2,$$

where Φ denotes the density of the standard normal distribution and $E = Up, Middle$. The overall likelihood is the product of l over all subjects, information conditions and choice lists.

[6] Under subjective expected utility $W_0(Up) = P_0(Up)$.

5. Results

5.1. Consistency

Consistency was good. We observed no significant differences between original and repeated ask prices in the two tests that we performed and the correlations were substantial (0.86 and 0.81, both $p < 0.01$). The mean absolute differences between the ask prices were 1.09 and 1.00 in the two questions.

A comparison between the ask prices of the option 20_U0 for stocks 1 and 2 (see Table 1) gives another consistency test. In both questions the subjects had no information about the underlying stock and they might treat them similarly. We indeed found no differences between the elicited ask prices and the correlation, although lower than in the other consistency tests, was still high and clearly different from 0 ($\rho = 0.52$, $p < 0.01$). The mean absolute error was equal to 1.53.

5.2. Subjective expected utility

Appendix A shows the median ask prices under the three informational conditions. Under expected utility, the subjective probabilities of the events are $(\frac{CE_j}{20})^{\beta}$ with β the power coefficient obtained from the estimation of utility. Overall, there was little utility curvature, which is consistent with the hypothesis that utility is about linear for small stakes (Wakker 2010). The median power coefficient was equal to 1 (mean 1.41, interquartile range = [0.82, 1.25]) and the proportion of subjects with concave utility did not differ from the proportion of subjects with convex utility.

Under expected utility, the subjective probabilities should satisfy complementarity and binary additivity, as discussed in Section 3. Panel A of Figure 4 shows that complementarity held in general. We could not reject the hypothesis that *P(MiddleUp)* + *P(Down)* = 100% for all three informational conditions. Moreover, we could not reject the hypothesis that the proportion of subjects for whom the sum of *P(MiddleUp)* and *P(Down)* exceeded 100% and the proportion for whom this sum was less than 100% were the same.

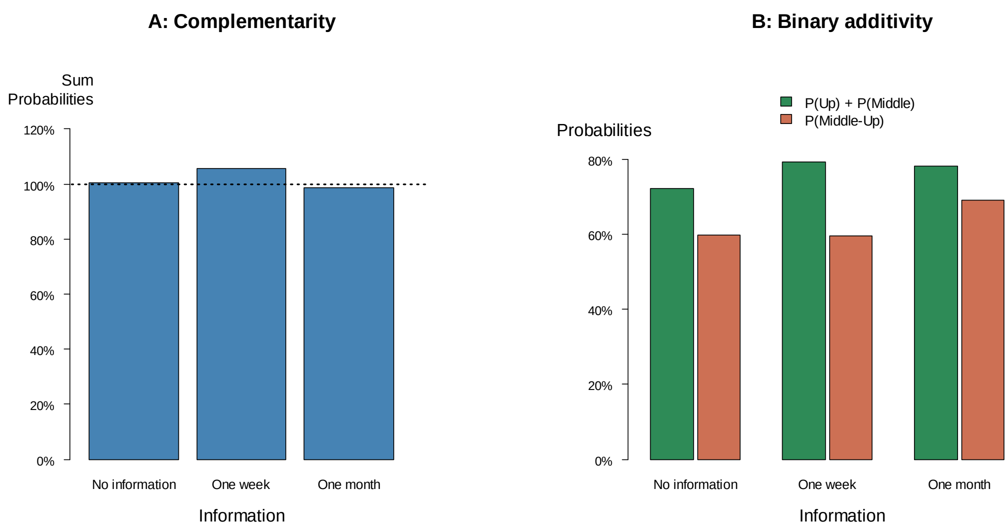

Figure 4. Tests of complementarity and binary additivity under subjective expected utility. Panel A shows that complementarity (*P(Down)* + *P(MiddleUp)* = 100%) held approximately. Panel B shows that binary additivity (*P(Up)* + *P(Middle)* = *P(MiddleUp)* was violated.

However, Panel B shows that binary additivity was violated. The sum of *P(Up)* and *P(Middle)* exceeded *P(MiddleUp)* in all three conditions suggesting binary subaditivity instead of binary additivity (all $p < 0.01$). The individual analyses confirmed this: the proportion of subjects who behaved according to binary subadditivity was significantly larger than the proportion of subjects displaying binary superadditivity in all three conditions (all $p < 0.01$).

Figure 4B also shows that the violations of binary additivity were smaller in the one month condition than in the other two conditions. Indeed, $P_t(Up) + P_t(Middle) - P_t(MiddleUp)$ was lower in the one month condition than in the no information condition ($p = 0.05$) and in the one week condition ($p < 0.01$). This shows that the deviations from expected utility decreased with more information.

The joint findings of complementarity and binary subadditivity are in line with previous evidence using introspective judgments (Tversky and Koehler 1994, Fox and Tversky 1998, Kilka and Weber 2001) and choice (Baillon and Bleichrodt 2012). They are consistent with support theory, a psychological theory of the formation of subjective probabilities (Tversky and Koehler 1994).

In summary, our data showed that subjective expected utility did not hold. There was some evidence that the violations decreased with more information in the one month condition.

5.3. Neo-additive model

Subjects whose subjective probabilities were outside the unit interval did not satisfy the neo-additive model and were excluded from the analyses. We only excluded these subjects for the informational condition for which they violated the neo-additive model, but not for the other conditions. This left 54 subjects in the no information condition, 50 subjects in the one week condition, and 46 subjects in the one month condition.

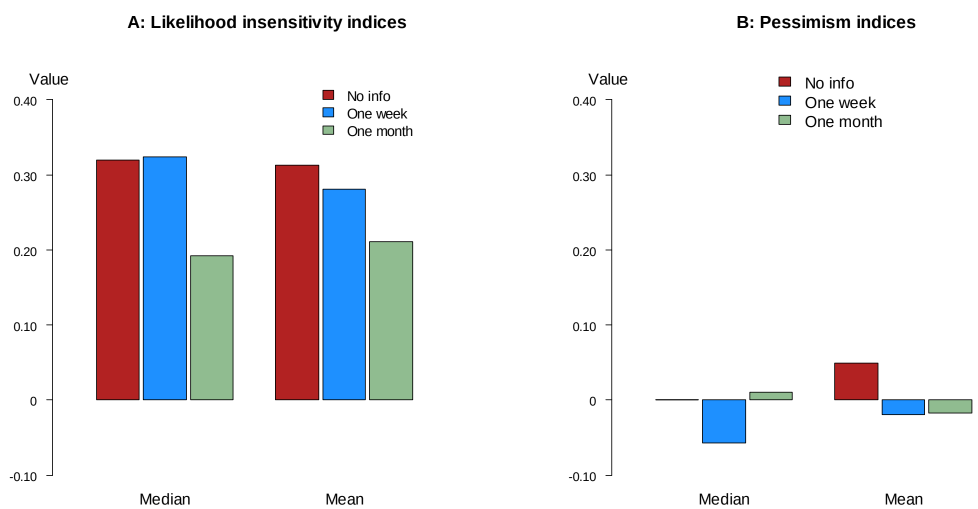

Figure 5. The likelihood insensitivity and pessimism indices. Panels A and B show the medians and means of the likelihood insensitivity and pessimism indices for the three informational conditions. Likelihood insensitivity falls with more information, but information has no effect on the pessimism indices.

Likelihood insensitivity

Figure 5A shows the likelihood insensitivity indices for the three conditions. All indices differed from zero ($p < 0.01$ in all three tests). They decreased with more information: the median decreased from 0.32 in the no information condition to 0.19 in the one month condition and the mean decreased from 0.31 to 0.21. The likelihood insensitivity index for one month was smaller than the index for no information ($p = 0.04$) and marginally smaller than the index for one week ($p = 0.08$). The indices for no information and for one week did not differ.

Figure 6 displays the relations between the individual values of the likelihood insensitivity (LIS) indices for the three informational conditions. In each of the panels, the horizontal axis shows the condition in which less information was available. The likelihood insensitivity indices were equal for points on the diagonal. If likelihood insensitivity decreased with the amount of information, then data points should be located below the diagonal. This is not so in Panel A, which compares the no information and the one week condition, but it is true in Panels B and C, confirming that likelihood insensitivity was lower in the one month condition than in the other two informational conditions.

Figure 6 also shows that a few subjects had negative likelihood insensitivity indices and were too sensitive to likelihood information. For these subjects, oversensitivity tended to decrease with information.

Overall, subjects moved in the direction of “correct” sensitivity to likelihood when they received more information. The shaded areas of Figure 6 show the subjects who moved in the direction of correct sensitivity: their likelihood insenstitivity or oversensitity decreased but they did not overshoot and went from insensitivity to even larger oversensitivity or from oversensitivity to even larger insensitivity. While many subjects are located outside the shaded area in Panel A, in Panels B and C most points are in the shaded area consistent with convergence towards expected utility.

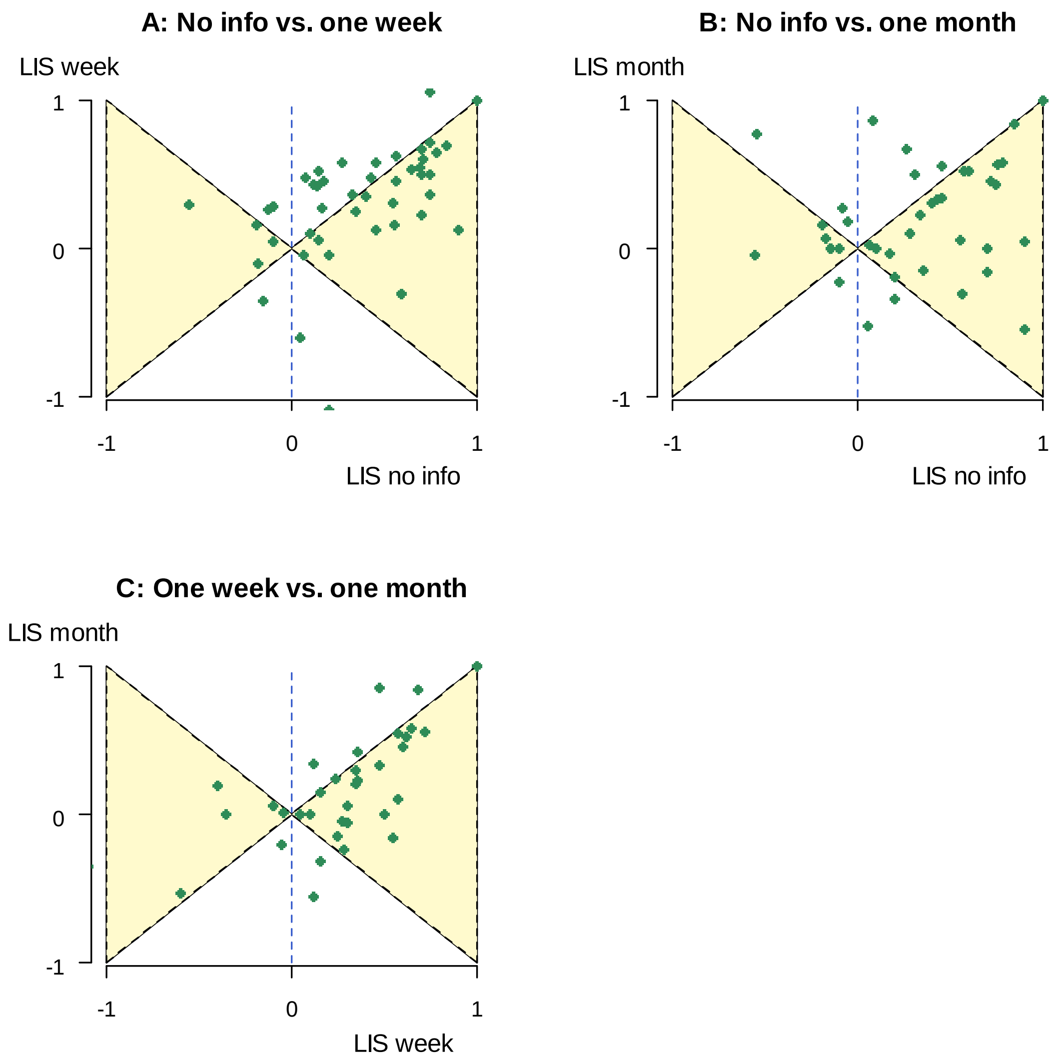

Figure 6. The relations between the individual likelihood insensitivity (a_t) indices. If subjects converge to expected utility then the points should lie in the shaded areas. This is so in Panels B and C.

The correlations between the likelihood insensitivity indices were fair to moderate: the Spearman correlation was 0.69 between no information and one week, 0.28 between no information and one month, and 0.58 between one week and one month.

Pessimism

Figure 5B shows the median and mean values of the pessimism indices for the three informational conditions. The median values were slightly negative, indicating optimism. The pessimism indices did not differ from zero. Figure 5B shows more optimism for the one week condition, but only the difference between the no information and the one week conditions was significant ($p < 0.01$).

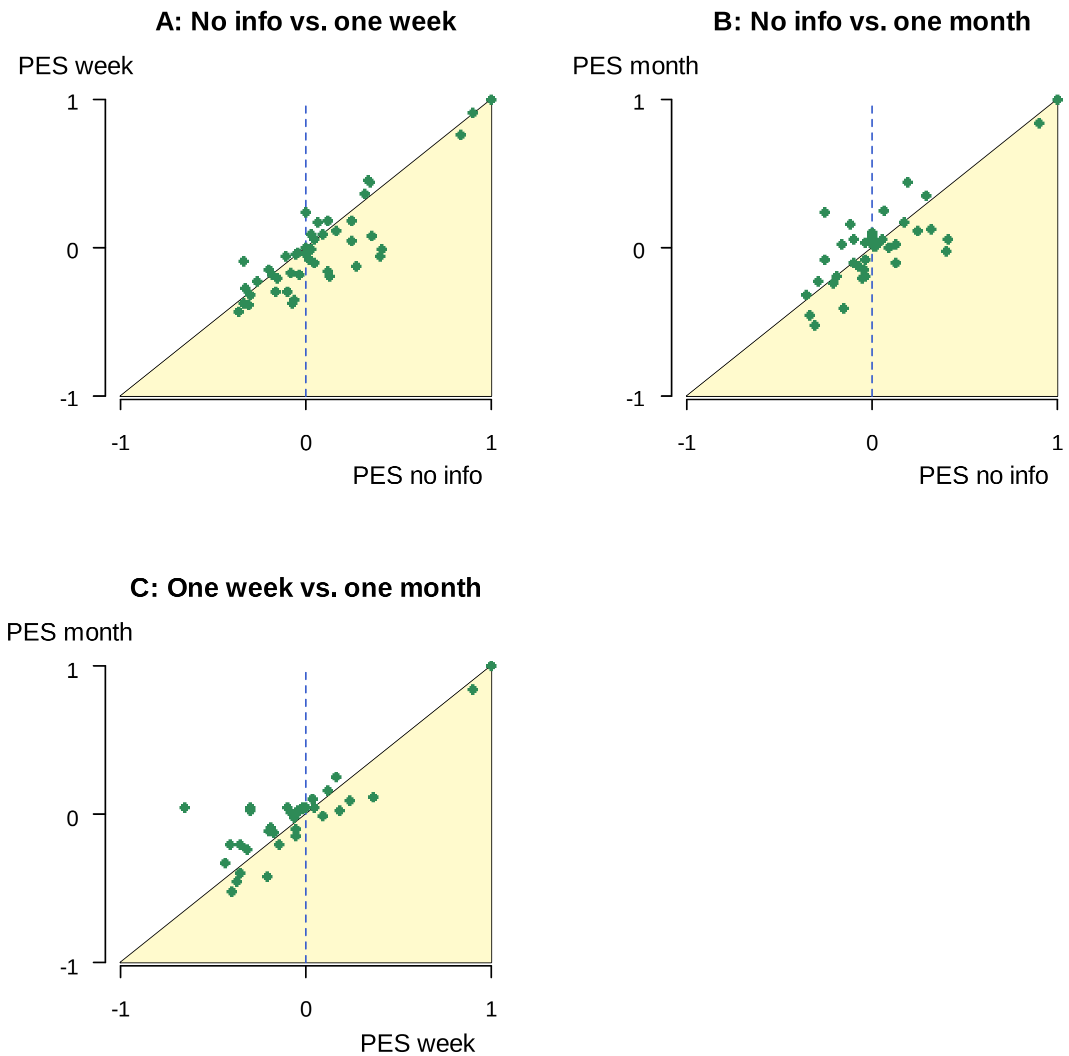

Figure 7. The relations between the individual pessimism (b_t) indices. If pessimism falls with more information then the points should be located in the shaded areas.

Figure 7 plots the individual pessimism (b_t) indices for the three informational conditions. In each panel, the condition with less information is plotted on the horizontal axis. Points on the diagonal have the same pessimism indices.

If pessimism decreases with information then individual points should be in the shaded lower halves of the figures. There was no evidence for this. If differences existed then they were small. The figure suggests that pessimism was a stable trait of individuals as the points were clustered around the diagonal. The correlations between the pessimism indices were substantial. The Spearman correlation was 0.79 between the no information and the one week conditions, 0.68 between the no information and one month conditions, and 0.77 between the one week and the one month conditions. They were also higher than the correlations between the likelihood insensitivity indices. Likelihood insensitivity was less stable than pessimism and it was affected more by new information.

5.4. Stochastic preferences

Table 2 summarizes the results of the maximum likelihood estimation. We estimated two models. The first model used the same data as the previous analysis and excluded choices that violated monotonicity. The second model used all individual choices, including those that violated monotonicity.

The results of the maximum likelihood estimation confirmed the individual analyses. Likelihood insensitivity was similar in the no information and the one week conditions, but it decreased in the one month condition. We found this in both models, but it was stronger in model 1. Our subjects were slightly optimistic (ambiguity seeking), particularly in the one week condition.

Subjects converged to expected utility with more information, but the null hypothesis that the parameters for the one month condition were equal to the expected utility parameters could be rejected (all $p < 0.01$). Both the likelihood insensitivity and the pessimism index differed from zero.

Table 2: Maximum likelihood estimations

		Model 1	Model 2
Violations of monotonicity included		No	Yes
a (likelihood insensititivity)		0.353***	0.384***
		(0.032)	(0.033)
	(one week effect)	−0.058	−0.032
		(0.046)	(0.047)
	(one month effect)	−0.172***	−0.089*
		(0.047)	(0.048)
b (pessimism)		−0.078***	−0.079***
		(0.018)	(0.019)
	(one week effect)	−0.064***	−0.052***
		(0.015)	(0.016)
	(one month effect)	−0.018	−0.004
		(0.016)	(0.016)
P(Up)	No info	0.394***	0.378***
		(0.01)	(0.011)
	1 week	0.334***	0.355***
		(0.009)	(0.011)
	1 month	0.333***	0.333***
		(0.008)	(0.01)
P(Middle)	No info	0.269***	0.283***
		(0.009)	(0.01)
	1 week	0.289***	0.278***
		(0.009)	(0.01)
	1 month	0.357***	0.365***
		(0.009)	(0.011)
Utility		0.899***	0.888***
		(0.022)	(0.022)
Noise		0.097***	0.098***
		(0.003)	(0.003)
Tremble		0.047***	0.073***
		(0.004)	(0.005)
Log-likelihood		−6550.2	−7661
N		128640	138240

Standard errors in parentheses.
***: significant at 1%, **, significant at 5%, * significant at 10%.

The multiple-priors interpretation

As explained in Section 3, we can also interpret our results in a multiple priors setting. Figure 8 shows the size of the set of priors for the three informational conditions using model 1 in Table 2. The black dot shows the estimated beliefs P_t. Together with a_t these determine the set of priors (the light grey area). The set of priors decreases with more information and is smallest in the one month condition.

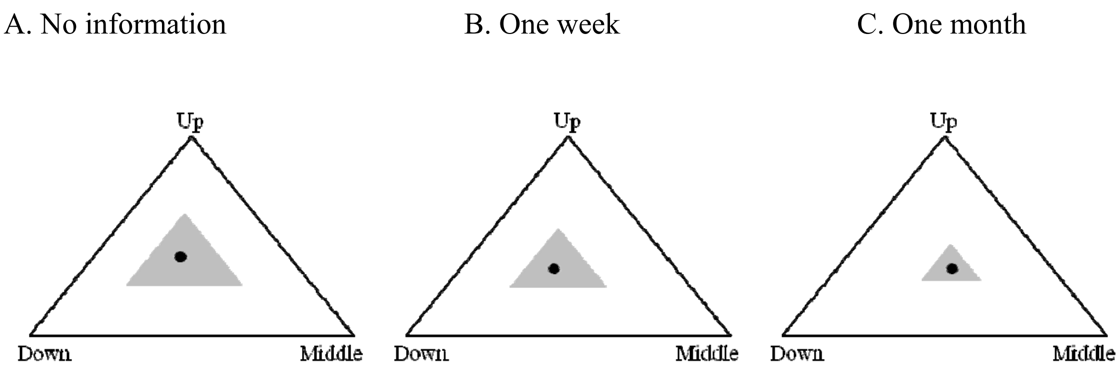

Figure 8: Sets of priors for the three informational conditions based on the estimates of Model 1. In each panel, the large triangle is the simplex representing all possible probability measures over the 3 events Up, Down, and Middle. Each vertex of the simplex denotes an event and corresponds to the measure in which this event is certain. Each opposite side of a vertex represents the probability measures assigning zero probability to the vertex event. The grey triangle is the set of priors and the black dot represents P_t.

Table 3: Maximum likelihood estimations – Multiple priors

		Model 1	Model 2
Violations of monotonicity included		No	Yes
$alpha_t$		0.389***	0.397***
		(0.030)	(0.028)
	(one week effect)	−0.130***	−0.083**
		(0.048)	(0.037)
	(one month effect)	−0.156**	−0.037
		(0.080)	(0.039)

Standard errors in parentheses. ***: significant At 1%, **, significant at 5%, * significant at 10%.

Table 3 shows the maximum likelihood estimates for $alpha_t$. In both models $alpha_t$ was less than 0.50 in the no information condition, which means that more weight was given to the highest expected utility. The representative subject was optimistic. This optimism increased in the one week condition and also in the one month condition. However, the two models disagree whether the increase in optimism was significant in the one month condition and whether it was larger than in the one week condition.

The theoretical literature distinguishes three methods for updating decision weights (Eichberger et al. 2010, Gilboa and Marinacci forthcoming): optimistic updating (Gilboa and Schmeidler 1993), the Dempster-Shafer rule which implies pessimistic updating (Dempster 1968, Shafer 1976), and full or generalized Bayesian updating (Walley 1991, Jaffray 1992, Chateauneuf et al. 2011). Eichberger et al. (2010) derived how the parameters of the neo additive model are updated under each of these three rules. Optimistic updating predicts that $alpha_t$ will go to 0, whereas pessimistic updating predicts that it will go to 1. Full Bayesian updating predicts that $alpha_t$ will remain constant. Table 3 shows that the results of Model 1 support optimistic updating, but we could not reject full Bayesian updating based on the results of Model 2.

Subjective probabilities

Table 2 also shows subjects' beliefs about the different events. P(Up) and P(Down) tend to decrease with more information, whereas P(Middle) increases.[7] The elicited probabilities were well-calibrated. For each day from their introduction to the 21st trading day we computed the proportions of the 328 IPOs that went up by more than 0.5% (corresponding with the event Up), that varied by at most 0.5% (corresponding with the event Middle), and

[7] Most differences are significant $p < 0.01$) except for the differences between P(Up) in the one week and the one month condition, between P(Middle) in the one week and the no information condition, and between P(Down) in the no information and the one month condition.

that went down by more than 0.5% (corresponding with the event Down). A frequentist may interpret these proportions as the actual probabilities of the events Up, Middle, and Down at each date t in the history.

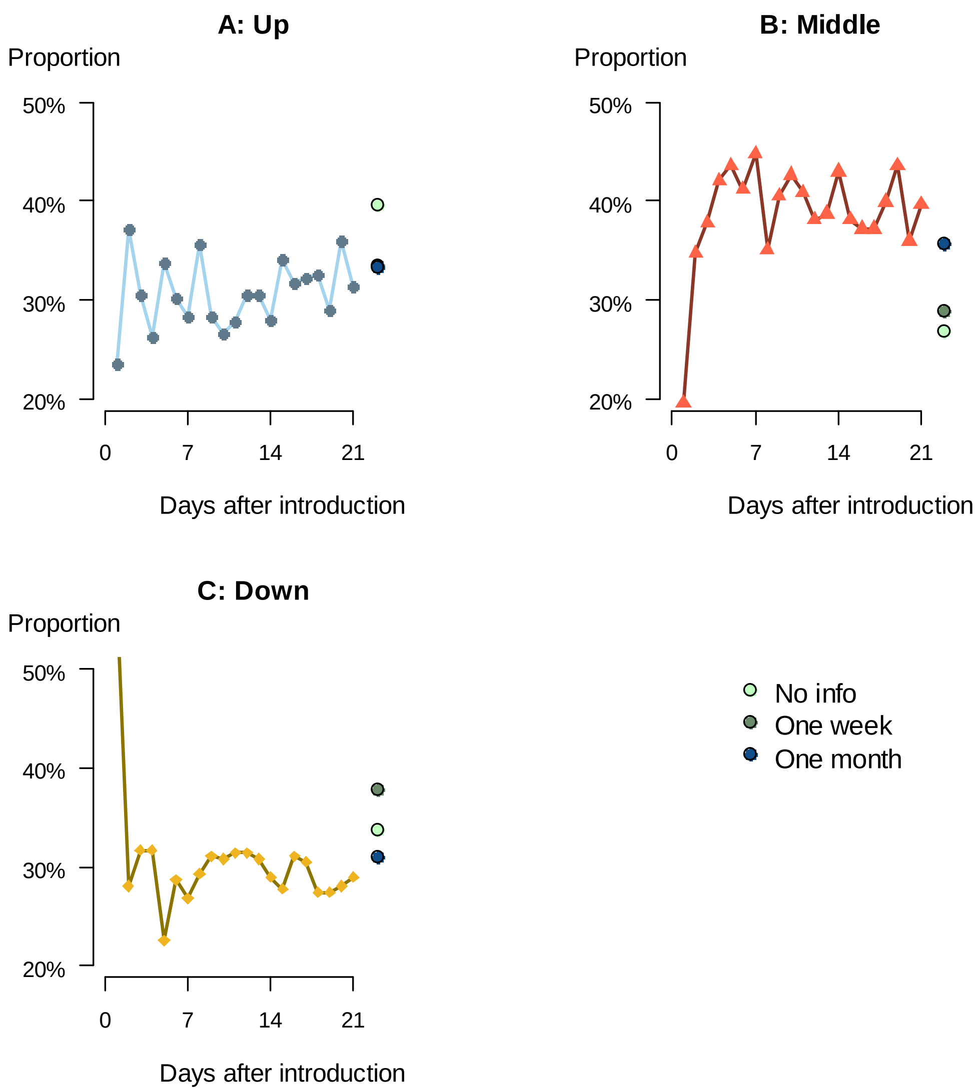

Figure 9. Stock history and beliefs. Panel A shows the proportion of the 328 IPOs that went up by more than 0.5% on each trading day from their introduction to the 21st trading day. Panels B and C show the proportions that varied by at most 0.5% and went down by more than 0.5%, respectively. The dots at the end show the estimated probabilities of *P(Up)* (Panel A), *P(Middle)* (Panel B), and *P(Down)* (Panel C) under the three informational conditions (in Panel A the points for one week and one month overlap).

Figure 9 shows the results of this analysis. Panel A shows the proportions for the event Up, Panel B for the event Middle, and Panel C for the event Down. The figure also shows the estimated median probabilities of *P(Up)*, *P(Middle)*, and *P(Down)* for the three informational conditions. They are displayed as dots at the end of the line.

All subjective probabilities converged to the actual probabilities. Subjects initially overestimated the probabilities of the events Up and Down. With more information, they adjusted their estimates downwards although a slight overestimation remained. On the other hand, subjects underestimated the probability of the event Middle. This underestimation decreased with information, particularly in the one month condition, but some remained.

In summary, Figure 9 shows that the representative subject overestimated the frequency of the more extreme events Up and Down and underestimated the frequency of the intermediate event Middle. The over- and underestimation was least in the one month condition and subjects correctly adjusted their probabilities when they received new information. The poor adjustment in the one week condition may be because the returns of IPOs are highly volatile in the first few trading days and the information provided was therefore not really informative. The final probabilities were well-calibrated and close to the aggregate frequencies in the market.

6. Discussion

This paper has studied the effect of learning new information on ambiguity attitudes using a simple method to correct beliefs for likelihood insensitivity and pessimism. Likelihood insensitivity decreased with more information. In the no information and the one week conditions, we found big likelihood insensitivity, even though we used experienced subjects. In the one month condition, likelihood insensitivity fell by half. Subjects went in the direction of correct sensitivity to likelihood information and converged to expected utility.

Likelihood insensitivity is often seen as a cognitive bias. Our findings suggest that this cognitive bias is reduced with more information. The conclusion that more information reduced cognitive biases also held under the multiple priors interpretation where information reduced the decision maker's set of priors.

Information had no effect on pessimism and the correlations between the pessimism indices were high for the three informational conditions. Pessimism may be a stable trait of decision makers that is unaffected by new information. This finding is consistent with the suggestion that pessimism reflects the motivational part of ambiguity attitude (Wakker 2010). If people are inclined to be pessimistic, then new information does not change this inclination.

We found little pessimism and in the one week condition even some optimism. In the multiple priors analysis we found significant optimism at the aggregate level. Our experiment used stock options and our subjects were finance students who knew about stocks. Empirical evidence suggests that ambiguity aversion decreases when subjects feel competent about the source of uncertainty and this may have explained our findings of no or little pessimism.

Another reason for the low amount of pessimism might be the use of ask prices in the elicitation of the certainty equivalents. Ask prices can lead to endowment effects (see e.g. Kahneman, Knetsch, and Thaler 1991) and, consequently, to an overestimation of certainty equivalents. This would lead to more optimism (see also Trautmann et al. 2011). The effect of endowment effects is the same for the three informational conditions similarly and, hence, they do not affect our conclusions about the effect of learning on ambiguity attitudes and beliefs.

The joint findings of close-to-zero pessimism and of decreasing likelihood insensitivity with more information suggest convergence towards expected utility with more information. This agrees with previous findings that experience and learning reduce biases.

On the other hand, the likelihood insensitivity index differed significantly from zero even in the one month condition. Moreover, under expected utility, the subjective probabilities violated binary additivity. Even though information moved our subjects in the direction of expected utility, important deviations remained.

Utility curvature had little impact and the results were similar when we assumed linear utility. If utility is linear, our method makes it possible to compute beliefs and ambiguity attitudes directly from the raw data without the need for any econometric estimation. This is an advantage for practical applications that seek to correct beliefs for ambiguity attitudes and cognitive biases. The similarity in results also suggests that utility curvature played only a minor role in decisions under uncertainty.

Our results were robust to the method of analysis. In the paper we reported two methods, one deterministic but allowing for some preference imprecision, the other stochastic. We also tried other approaches, including interval arithmetic computation to account for imprecision in preferences. The results of these additional analyses were similar.

We assumed that utility did not depend on the information about past events. The utility function reflects preferences over outcomes and new information about the state space has no relevance for these. Abdellaoui et al. (2011) measured utility for different sources of uncertainty and could not reject the null hypothesis that utility was the same across sources.

A more controversial assumption is that probabilistic sophistication holds within histories and, hence, that subjective probabilities exist. Different histories can be interpreted as different sources of uncertainty. The notion of sources of uncertainty was first proposed by Amos Tversky in the 1990s (Tversky and Kahneman 1992, Tversky and Fox 1995, Tversky and Wakker 1995). Chew and Sagi (2006, 2008) showed that, if an exchangeability condition holds, subjective probabilities can be defined within sources even when probabilistic sophistication does not hold between sources. Our analysis implicitly assumed this condition.

Abdellaoui et al. (2011) obtained support for it in all but one of their tests. The only exception was a test involving an unfamiliar source and hypothetical choice. For real incentives, exchangeability always held. The real incentive system they used was similar to ours. Moreover, because our subjects were finance students, all sources were familiar. Further, the estimated beliefs were well-calibrated: they were sensitive to new information and they reflected aggregate market behavior.

We finally assumed that the weighting function could be described by the neo-additive form. This assumption is not too restrictive as the neo-additive weighting function provides a good approximation to more general weighting functions (Diecidue et al. 2009, Abdellaoui et al. 2010). For most subjects the estimated model parameters were plausible and within the range allowed by the model.

7. Conclusion

Ambiguity theories are useful for studying the effects of learning in decision under uncertainty. Learning affects both beliefs and ambiguity attitudes. We developed a new method that corrects beliefs for ambiguity attitudes. Our method decomposes ambiguity attitudes into likelihood insensitivity and pessimism. Likelihood insensitivity decreased with new information, but pessimism was unaffected. This is consistent with the interpretation that likelihood insensitivity is a cognitive bias and that pessimism is the motivational part of ambiguity attitude. Subjects behaved more in line with expected utility when they received more information, but significant deviations from expected utility remained. The estimated beliefs, when corrected for ambiguity attitudes, moved in the direction of aggregate market behavior indicating that subjects took the information into account and reacted to it in a reasonable manner.

Appendix A: Median ask prices

Option	Up	Middle	Down	Middle-Up
No information	8.50	7.50	7.50	12.50
1 week	8.50	8.00	9.00	12.50
1 month	7.50	7.50	7.50	13.50

Appendix B (Experimental instructions)

Instructions

Thank you for participating in our experiment. For your participation, you will receive a show up fee of €5 and an extra payment depending on your choices during the experiment. Please read the instructions carefully. Before starting the experiment, we will ask you several questions to test your understanding of the instructions. If you answer every question correctly, you will proceed to the experiment; otherwise, we will ask you to read the instructions once more and re-answer the questions until all your answers are correct. We want to be sure that you have understood the instructions so that your answers in the experiment reflect your preferences and are not caused by any misunderstandings. If you have any questions, please feel free to ask the experimenter.

During the experiment, you have to answer a series of choice questions. There are **no right or wrong answers** to these questions. We are interested in ***your*** preferences. Your final payment will be determined by the choices you make during the experiment. Hence it is in your own interest to reveal your true preferences in the choices you will face.

***

During the experiment, you will be asked to choose between a ***digital option*** for an underlying stock and a sure money amount.

A digital option for an underlying stock pays a pre-specified money amount **H** if a given event occurs and **L** otherwise.

The **underlying stock** is randomly chosen from a database of stocks that were newly-listed on the NYSE between 1 **January** 2009 and 25 February 2011. The stocks in the database are randomly numbered from 1 to 328. At the beginning of the experiment, you will draw 4 numbers from a box, and the 4 corresponding stocks will be used as the underlying stocks of your digital options. At the end of the experiment, the names of the stocks will be revealed, and you can check the historical quotes of the stock prices on Yahoo Finance afterwards. Note that we cannot manipulate the price distribution of the stocks as these are historically given.

You will face 3 different situations.

Situation 1

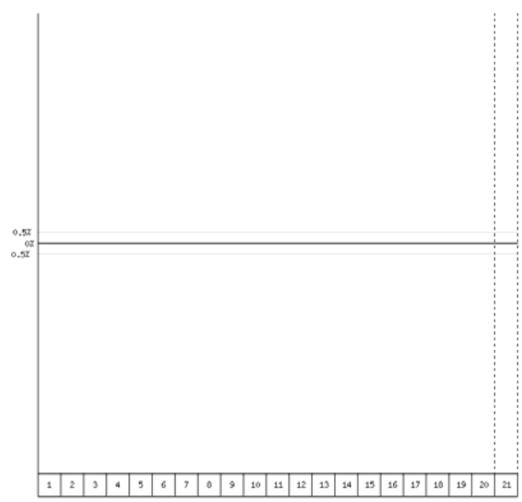

- **Situation 1:** You have an option for an underlying stock, which has just been listed on the Stock Exchange. Consequently, you have no quotes of the historical stock price. You know that the expiration date of the option is the 21st trading day of the stock, and the payoff of the option depends on the daily return of the stock on the 21st trading day. (More explanation about the option payoff will be presented later.)

Situation 2

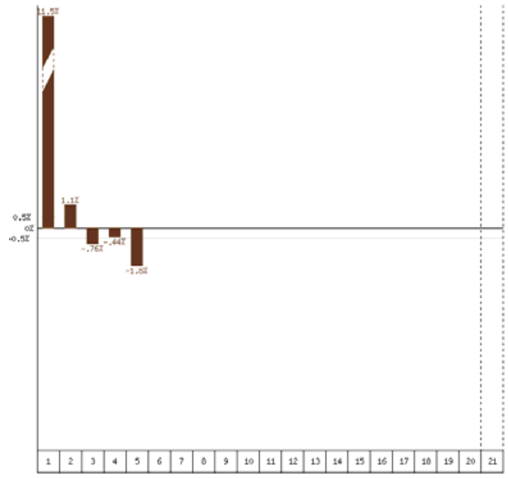

- **Situation 2:** You have an option for an underlying stock, which has been listed on the Stock Exchange for one week. You have 5 quotes of the historical daily return of the stock, which have been depicted by the brown bars. You know that the expiration date of the option is the (same) 21st trading day of the stock, and the payoff of the option depends on the daily return of the stock on the 21st trading day.

Situation 3

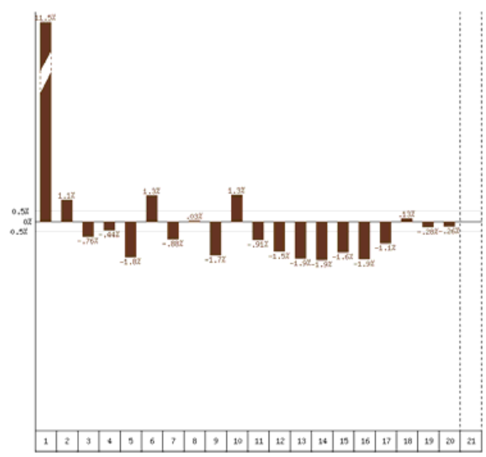

- **Situation 3:** You have an option of an underlying stock, which has been listed on the Stock Exchange for 20 days. You have 20 quotes of the historical daily return of the stock, which have been depicted by the brown bars. You know that the expiration date of the option is the (same) 21st trading day of the stock, and the payoff of the option depends on the daily return of the stock on the 21st trading day.

You will face 4 types of *digital option*

For each situation described above, you may face 4 types of digital options. Here, we use the first situation as an example to illustrate the 4 types of digital options.

Up-Option

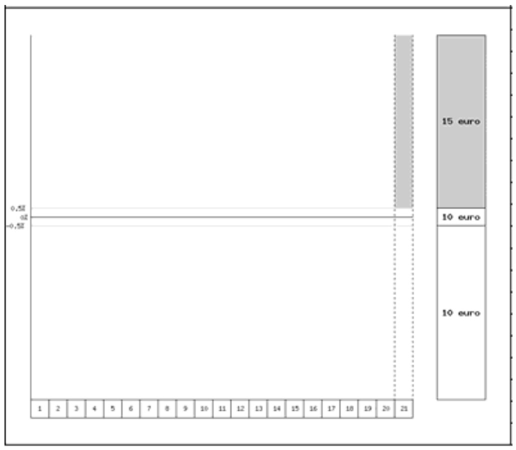

Middle-Option

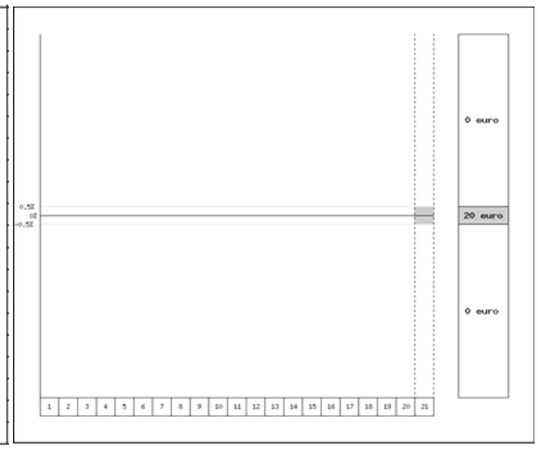

Down-Option

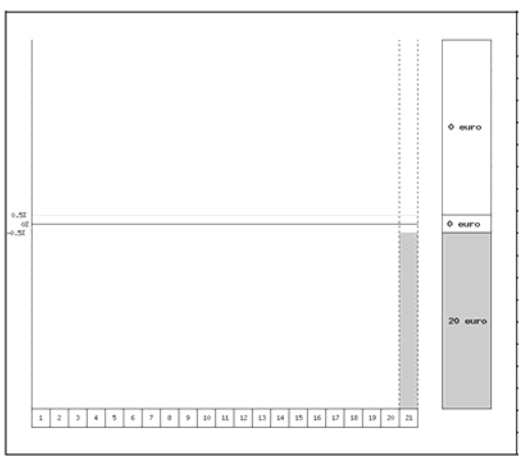

MiddleUp-Option

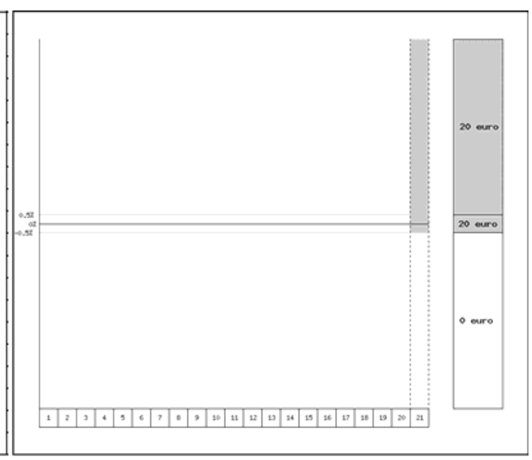

- An **Up-option** pays €H if the daily return (r) of the underlying stock on its expiration day exceeds +0.5% ($r > +0.5\%$) and €L otherwise;
- **A Middle-option** pays €H if the daily return (r) of the underlying stock on its expiration day varied between –0.5% and +0.5% ($-0.5\% \leq r \leq +0.5\%$) and €L otherwise;
- A **Down-option** pays €H if the daily return (r) of the underlying stock on its expiration day is less than –0.5% ($r < -0.5\%$) and €L otherwise.
- A **MiddleUp-option** pays €H if the daily return (r) of the underlying stock on its expiration day exceeds –0.5% ($r \geq -0.5\%$) and €L otherwise;

€H and €L are pre-specified money amounts. For instance, the figure above displays an Up-option with H=15 and L=10, and the other three types with H=20 and L=0. You may encounter different H and L in the experiment.

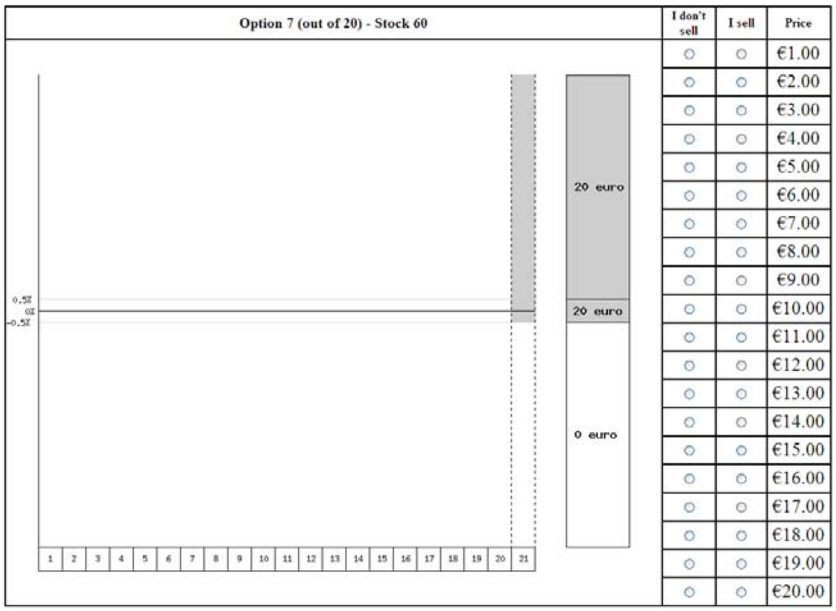

We will determine your selling price of 20 different options through a series of choices between **the option** and **a certain money amount**. An example is given in the above figure. For each of the 20 prices, you are asked to indicate whether you would like to sell the option or not. The money amount where you switch your choice from 'I don't sell' to 'I sell' is taken as your selling price. All sales will be realized on the 21^{st} day.

If you sell at €x, do you agree that you also want to sell at prices higher than €x? Y/N

If you don't sell at €y, do you agree that you don't want to sell at prices lower than €y? Y/N

Payment

After making all the 20 choices, please call the experimenter. The experimenter will let you throw a 20-sided dice twice. The first throw will determine which of your choices will be played for real. The second throw determines the price you are offered.

As an example, imagine that you throw 7 on your first throw and 6 on your second. Hence the 7^{th} choice will be selected and the price you are offered for the option in the 7^{th} choice is €6.

Suppose that option in the 7th choice is a MiddleUp-option with H=20 and L=0, as in the figure above. Suppose further that your selling price for the 7th option was found to be €9. This means that you are not willing to sell the option for a price less than €9 and, hence, you don't accept the offered price of €6 and **thus you keep the option**;

- If the daily return on the 21st trading day of the underlying stock is at least –0.5% (e.g. 0.15%), then we pay you €20 plus the €5 show-up fee. In total you get €25.
- If the daily return on the 21st trading day of the underlying stock is smaller than –0.5% (e.g. –1.49%), then we pay you €0 plus the €5 show-up fee. In total you get €5.

Now imagine that you throw 7 and 10. Then the price offer you are offered is €10. Because you are willing to sell the option if the price is at least €9, you accept the offered price of €10 and thus we pay you €10 plus the €5 show-up fee. In total you get €15.

Note that it is in your best interests to state your selling price truthfully. To see that, suppose your true selling price is €9, but you state a selling price of €11. Then if the price we offer for the option is €10, you keep the option even though it is worth less to you than €10.

Questions:

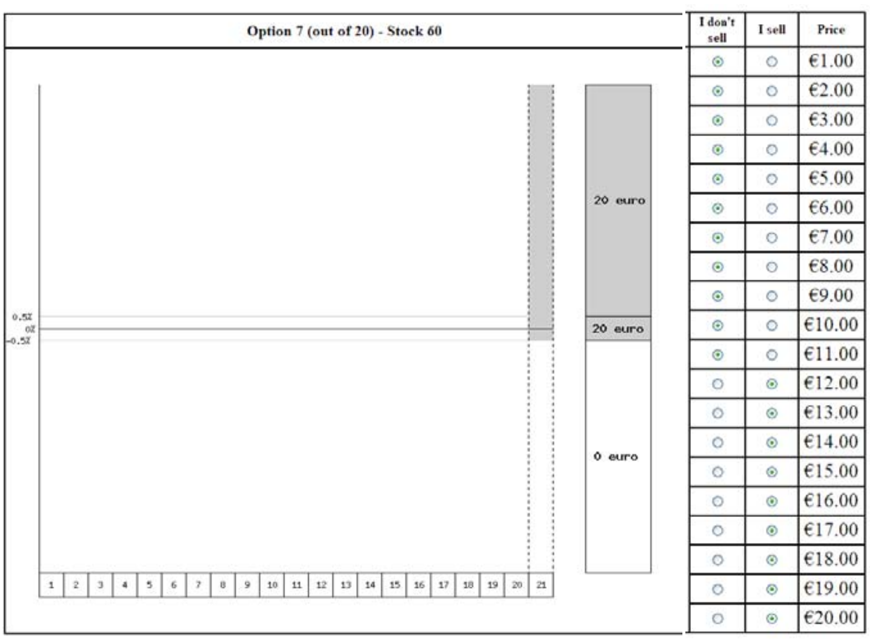

Suppose you are going to play the choice in the picture above for real.

1. What is the minimum selling price?
2. What is the payoff of the plotted option, if the daily return on the 21^{st} trading day is:
   - 1.4%?
   - –0.45%?
   - –1.4%?
3. Suppose that the daily return on the 21^{st} trading day is 1.4%, what is the total payment you get if the second number you throw is 1?
4. Suppose the daily return on the 21^{st} trading day is 1.4%, what is the total payment you get if the second number you throw is 15?

References

Abdellaoui, M., A. Baillon, L. Placido, P. P. Wakker. 2011. The rich domain of uncertainty: Source functions and their experimental implementation. *American Economic Review* **101** 695-723.

Abdellaoui, M., O. l'Haridon, H. Zank. 2010. Separating curvature and elevation: A parametric probability weighting function. *Journal of Risk and Uncertainty* **41** 39-65.

Abdellaoui, M., F. Vossmann, M. Weber. 2005. Choice-based elicitation and decomposition of decision weights for gains and losses under uncertainty. *Management Science* **51** 1384-1399.

Abdellaoui, M., H. Bleichrodt, O. l'Haridon. 2008. A tractable method to measure utility and loss aversion under prospect theory. *Journal of Risk and Uncertainty* **36** 245-266.

Baillon,A. and H. Bleichrodt. . 2012. Testing ambiguity models through the measurement of probabilities for gains and losses. *Working Paper Erasmus School of Economics.*

Bleichrodt, H., J. L. Pinto, P. P. Wakker. 2001. Making descriptive use of prospect theory to improve the prescriptive use of expected utility. *Management Science* **47** 1498-1514.

Bostic, R., R. J. Herrnstein, R. D. Luce. 1990. The effect on the preference reversal of using choice indifferences. *Journal of Economic Behavior and Organization* **13** 193-212.

Bruhin, A., H. Fehr-Duda, T. Epper. 2010. Risk and rationality: Uncovering heterogeneity in probability distortion. *Econometrica* **78** 1372-1412.

Charness, G., D. Levin. 2005. When optimal choices feel wrong: A laboratory study of bayesian updating, complexity, and affect. *American Economic Review* **88** 933-946.

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

Chateauneuf, A., T. Gajdos, J. Y. Jaffray. 2011. Regular updating. *Theory and Decision* **71** 111-128.

Chew, S. H., J. S. Sagi. 2006. Event exchangeability: Probabilistic sophistication without continuity or monotonicity. *Econometrica* **74** 771-786.

Chew, S. H., J. S. Sagi. 2008. Small worlds: Modeling attitudes toward sources of uncertainty. *Journal of Economic Theory* **139** 1-24.

Cohen, M., I. Gilboa, J. Jaffray, D. Schmeidler. 2000. An experimental study of updating ambiguous beliefs. *Risk, Decision, and Policy* **5** 123-133.

Conte, A., J. D. Hey, P. G. Moffatt. 2011. Mixture models of choice under risk. *Journal of Econometrics* **162** 79-88.

Delquié, P. 1997. 'Bi-matching': A new preference assessment method to reduce compatibility effects. *Management Science* **43** 640-658.

Dempster, A. P. 1968. A generalization of bayesian inference. *Journal of the Royal Statistical Society.Series B (Methodological)* **30** 205-247.

Di Mauro, C. 2008. Uncertainty aversion vs. competence: An experimental market study. *Theory and Decision* **64** 301-331.

Diecidue, E., U. Schmidt, H. Zank. 2009. Parametric weighting functions. *Journal of Economic Theory* **144** 1102-1118.

Dominiak, A., P. Dürsch, J. Lefort. 2012. A dynamic Ellsberg urn experiment. *Games and Economic Behavior* **75** 625-638.

Eichberger, J., S. Grant, D. Kelsey. 2007. Updating Choquet beliefs. *Journal of Mathematical Economics* **43** 888-899.

Eichberger, J., S. Grant, D. Kelsey. 2010. Comparing three ways to update Choquet beliefs. *Economics Letters* **107** 91-94.

Eichberger, J., S. Grant, D. Kelsey. 2012. When is ambiguity attitude constant? *Journal of Risk and Uncertainty* **45** 239-263.

El-Gamal, M. A., D. M. Grether. 1995. Are people bayesian? uncovering behavioral strategies. *Journal of the American Statistical Association* **90** 1137-1145.

Ellsberg, D. 1961. Risk, ambiguity and the Savage axioms. *Quarterly Journal of Economics* **75** 643-669.

Epstein, L. G. 2006. An axiomatic model of non-bayesian updating. *The Review of Economic Studies* **73** 413-436.

Epstein, L. G., M. Schneider. 2007. Learning under ambiguity. *The Review of Economic Studies* **74** 1275-1303.

Fox, C. R., M. Weber. 2002. Ambiguity aversion, comparative ignorance, and decision context. *Organizational behavior and human decision processes* **88** 476-498.

Fox, C. R., A. Tversky. 1998. A belief-based account of decision under uncertainty. *Management Science* **44** 879-895.

Gajdos, T., T. Hayashi, J. M. Tallon, J. C. Vergnaud. 2008. Attitude toward imprecise information. *Journal of Economic Theory* **140** 27-65.

Ghirardato, P., F. Maccheroni, M. Marinacci. 2004. Differentiating ambiguity and ambiguity attitude. *Journal of Economic Theory* **118** 133-173.

Ghirardato, P., M. Marinacci. 2001. Risk, ambiguity, and the separation of utility and beliefs. *Mathematics of Operations Research* **26** 864-890.

Gilboa, I., D. Schmeidler. 1993. Updating ambiguous beliefs. *Journal of Economic Theory* **59** 33-49.

Gilboa, I. & Marinacci, M. (forthcoming). Ambiguity and the bayesian paradigm. *Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society.* .

Gilboa, I., D. Schmeidler. 1989. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics* **18** 141-153.

Grether, D. M. 1980. Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly Journal of Economics* **95** 537-557.

Griffin, D., A. Tversky. 1992. The weighing of evidence and the determinants of confidence. *Cognitive psychology* **24** 411-435.

Hamermesh, D. S. 1985. Expectations, life expectancy, and economic behavior. *The Quarterly Journal of Economics* **100** 389-408.

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

Heath, C., A. Tversky. 1991. Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty* **4** 4-28.

Hey, J. D., C. Orme. 1994. Investigating generalizations of expected utility theory using experimental data. *Econometrica* **62** 1291-1326.

Hoffman, R. M., J. H. Kagel, D. Levin. 2011. Simultaneous versus sequential information processing. *Economics Letters* **112** 16-18.

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

Ju, N., J. Miao. 2012. Ambiguity, learning, and asset returns. *Econometrica* **80** 559-591.

Kahn, B. E., R. K. Sarin. 1988. Modeling ambiguity in decisions under uncertainty. *Journal of Consumer Research* **15** 265-272.

Kahneman, D., A. Tversky. 1972. Subjective probability: A judgment of representativeness. *Cognitive psychology* **3** 430-454.

Kahneman, D., A. Tversky. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* **47** 263-291.

Keynes, J. M. 1921. *A Treatise on Probability.* McMillan, London.

Kilka, M., M. Weber. 2001. What determines the shape of the probability weighting function under uncertainty? *Management Science* **47** 1712-1726.

List, J. A. 2004. Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica* **72** 615-625.

Luce, R. D. 1991. Rank-and sign-dependent linear utility models for binary gambles. *Journal of Economic Theory* **53** 75-100.

Machina, M. J., D. Schmeidler. 1992. A more robust definition of subjective probability. *Econometrica: Journal of the Econometric Society* 745-780.

McCord, M., R. de Neufville. 1986. Lottery equivalents: Reduction of the certainty effect problem in utility assessment. *Management Science* **32** 56-60.

Miyamoto, J. M. 1988. Generic utility theory: Measurement foundations and applications in multiattribute utility theory. *Journal of Mathematical Psychology* **32** 357-404.

Myagkov, M., C. R. Plott. 1997. Exchange economies and loss exposure: Experiments exploring prospect theory and competitive equilibria in market environments. *The American Economic Review* **87** 801-828.

Noussair, C., S. Robin, B. Ruffieux. 2004. Revealing consumers' willingness-to-pay: A comparison of the BDM mechanism and the Vickrey auction. *Journal of Economic Psychology* **25** 725-741.

Phillips, L. D., W. Edwards. 1966. Conservatism in a simple probability inference task. *Journal of Experimental Psychology; Journal of Experimental Psychology* **72** 346.

Poinas, F., J. Rosaz, B. Roussillon. 2012. Updating beliefs with imperfect signals: Experimental evidence. *Journal of Risk and Uncertainty* **44** 219-241.

Rabin, M., J. L. Schrag. 1999. First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics* **114** 37-82.

Schmeidler, D. 1989. Subjective probability and expected utility without additivity. *Econometrica* **57** 571-587.

Shafer, G. 1976. *A Mathematical Theory of Evidence.* Princeton University Press, Princeton, New Jersey.

Smith, V. K., D. H. Taylor Jr, F. A. Sloan. 2001. Longevity expectations and death: Can people predict their own demise? *American Economic Review* 1126-1134.

Stott, H. P. 2006. Cumulative prospect theory's functional menagerie. *Journal of Risk and Uncertainty* **32** 101-130.

Trautmann, S. T., F. M. Vieider, P. P. Wakker. 2011. Preference reversals for ambiguity aversion. *Management Science* **57** 1320-1333.

Tversky, A., D. Kahneman. 1973. Availability: A heuristic for judging frequency and probability. *Cognitive psychology* **5** 207-232.

Tversky, A., D. K. Koehler. 1994. Support theory: A nonexistential representation of subjective probability. *Psychological Review* **101** 547-567.

Tversky, A., C. Fox. 1995. Weighing risk and uncertainty. *Psychological Review* **102** 269-283.

Tversky, A., P. P. Wakker. 1995. Risk attitudes and decision weights. *Econometrica* **63** 297-323.

Tversky, A., D. Kahneman. 1992. Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* **5** 297-323.

van de Kuilen, G., P. P. Wakker. 2006. Learning in the Allais paradox. *Journal of Risk and Uncertainty* **33** 155-164.

Wakker, P. P. 2010. *Prospect Theory: For Risk and Ambiguity.* Cambridge University Press, Cambridge UK.

Wakker, P. P., D. Deneffe. 1996. Eliciting von neumann-morgenstern utilities when probabilities are distorted or unknown. *Management Science* **42** 1131-1150.

Walley, P. 1991. *Statistical Reasoning with Imprecise Probabilities.* Chapman and Hall, London.