Testing Negative Value of Information and Ambiguity Aversion*

Christopher Kops and Illia Pasichnichenko†

Abstract

The standard Subjective Expected Utility model of decision-making implies that information can never have a negative value ex-ante. Many ambiguity theories have since questioned this property. We provide an experimental test of the connection between the value of information and ambiguity attitude. Our results show that the value of information can indeed be negative when new information renders hedging against ambiguity impossible. Moreover, the value of information is correlated with ambiguity aversion. This confirms the predictions from ambiguity theories and may have implications for decision-making in uncertain and dynamic environments. Neither complexity avoidance nor information with ambiguous reliability can reproduce the results.

JEL codes: D81, D83, D91

Keywords: Value of Information; Information Aversion; Ambiguity Aversion; Ellsberg Urn; Dynamic Consistency; Complexity Avoidance

1 Introduction

Information is a means to resolve uncertainty. It has economic value. Examples run from the Subjective Expected Utility (SEU) model of decision-making under uncertainty combined with the standard updating rules (Blackwell, 1953; Savage, 1954) to games where information is covertly acquired (Pavan and Tirole, 2022). A lesser known characteristic of information is that it may increase perceived uncertainty, by turning a situation with known probabilities into

---

*This draft: August 23, 2023

†Kops: Maastricht University, The Netherlands (e-mail: j.kops@maastrichtuniversity.nl); Pasichnichenko: Queen Mary University of London, United Kingdom (email: i.pasichnichenko@qmul.ac.uk). We thank the editor Pierpaolo Battigalli, an associate editor, and two anonymous referees for the valuable comments and advice, as well as Aurélien Baillon, Peter Duersch, Jürgen Eichberger, Asen Ivanov, Jörg Oechsler, Marieke Pahlke, Hannes Rau, Frank Riedel, Arno Riedl, Stefan Trautmann, Christoph Vanberg, Joël van der Weele and audiences at Heidelberg, Bielefeld, Maastricht, Sussex, Bristol, Queen Mary, RUD 2021 (where this paper won the Jaffray Prize) and D-TEA 2022 for helpful comments and suggestions. An earlier version of this paper circulated under the title "A Test of Information Aversion". Part of this work was carried out while both authors were at Heidelberg University. Illia Pasichnichenko acknowledges financial support of the Alexander von Humboldt Foundation, Christopher Kops acknowledges financial support of the Graduate School of Business & Economics at Maastricht University.

one where probabilities are not known. This way, information can impede the decision-making process. It can even do so to a point where its value for a non-SEU decision maker (DM) becomes negative. Instances of this type, but unrelated to uncertainty, have been reported for decision-making by managers (Deshpande and Kohli, 1989), and investors (Karlsson, Loewenstein, and Seppi, 2009). This paper proposes a way to test the connection between the negative value of information and ambiguity aversion, and presents the results from an experimental implementation of that test.

Our design builds on the classic thought experiment of Ellsberg (1961) to test the SEU model. Specifically, our experiment features urns with balls where each ball is marked with one of the numbers 1, 2, 3, or 4. The distinctive feature of Ellsberg urns is that some information about their ball composition is missing. In our design, the numbers n_1 and n_2 of balls marked with 1 and 2 are known, whereas the numbers n_3 and n_4 are unknown. However, the total number $n_3 + n_4$ is known and is such that $n_1 + n_2 < n_3 + n_4$. Further, it is known that the balls marked with 1 and 3 are green, while the balls marked with 2 and 4 are blue (Figure 1).

Figure 1: Ball types

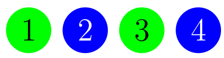

To test the negative value of information, subjects in our experiment are asked to choose between different bets on the number x a randomly drawn ball from the urn is marked with. They are asked to choose between the bet on 1 and 2 (win if $x \in \{1, 2\}$, lose otherwise) and the bet on 3 and 4 (win if $x \in \{3, 4\}$, lose otherwise). They are asked this question in two different information conditions. In one such condition, call it No Color Info, the color of the drawn ball is unknown. Note that in this condition, the winning probabilities of each bet are known. In the other condition, call it Color Info, they know the color of the drawn ball up front. Here, regardless of whether the drawn ball is green or blue, the winning probabilities become unknown. For example, if the drawn ball is green, subjects may want to compare n_1 and n_3, where n_3 is unknown. Our test of the negative value of information boils down to the question of which of the two conditions, No Color Info or Color Info, subjects prefer. By design, the payoffs in the Color Info condition are higher than in No Color Info. Thus, somebody who is indifferent to the information would prefer the Color Info condition, while a preference for No Color Info implies a negative value of the information.

Our results in this paper show that for over 49% of the participants in our experiments the value of information is negative. As laid out above, the intuition behind these experiments is that information breaks an existing hedge against ambiguity which may induce ambiguity averse subjects to avoid it. The data confirms this intuition. It shows a significant correlation between the negative value of information and ambiguity aversion. Further treatments establish that the negative value of information is neither driven by the complexity of the decision situation, nor can it be replicated by information of ambiguous reliability. This underlines the robustness of our results and corroborates the theoretical predictions on the relationship between the value of information and ambiguity attitude (Li, 2020).

As an application, consider the following example where all news is bad news. The stock market can be in one of two states, either up or down. An investor chooses between two investments. The first is a real estate investment which is typically not correlated with the stock market. The second concerns an asset where data on its correlation with the market is missing. Not knowing whether the market is up or down, the asset is more likely to be successful than the real estate investment is. Knowing the state of the market, however, makes the situation ambiguous, because the correlation between the market and the asset is unknown. This makes it conceivable that the investor ascribes a negative value to the information about the state of the market. It is in this way that the negative value of information may explain underinvestment.

The example above involves a variable (state of the stock market) which the DM deems relevant and whose relationship with other variables is ambiguous. Another way to induce a negative value of information lies in a gradual resolution of uncertainty. This way, a DM's perceived uncertainty increases, because partial information is revealed at intermediate stages (Li, 2020).[1]

Theoretically, information can never have a negative value ex-ante under the assumption of expected utility maximization and the standard updating rules (Blackwell, 1953). The idea that a non-expected utility theory may lead to a negative value of information was first presented in Wakker (1988) for decision making under risk. Li (2020) studies preferences that depend on how information is dynamically revealed. Specifically, the preference relation is defined on pairs (f, π), where f is an act and π is a partition of the state space representing anticipated information. The paper characterizes a recursive utility representation by a set of axioms which includes weak order and strong monotonicity. Different types of static ambiguity averse preferences can be embedded into the model. Our results in this paper may be seen as an experimental test for the axioms in Li (2020). With close to 80% of ambiguity averse subjects showing a negative value of information, our findings by and large confirm the theoretical prediction of Li (2020). On the other hand, we did not find significant support for the connection between the value of information and dynamic consistency in Galanis (2021), although the effect goes in the predicted direction. Future research may shed further light on these issues.

While we study the value of reliable information in this paper, a recent strand of the experimental literature studies the value of ambiguous information which is information with unknown or ambiguous reliability (Kellner, Le Quement, and Riener, 2022; Shishkin and Ortoleva, 2023), or unknown informativeness (Epstein and Halevy, 2020; Liang, 2021). Other studies tested consequentialism and dynamic consistency under uncertainty (Cohen et al., 2000; Dominiak,

[1] Genetic testing (Oster, Shoulson, and Dorsey, 2013) and checking investment accounts (Karlsson, Loewenstein, and Seppi, 2009) are typical examples here.

Duersch, and Lefort, 2012; Bleichrodt et al., 2021; Esponda and Vespa, 2021). Out of all these, our experiment is presumably closest to Shishkin and Ortoleva (2023). They, too, study the value of information and its relation to ambiguity attitude when 'all news is bad news' as termed by Gul and Pesendorfer (2021). However, they find no correlation between the negative value of information and ambiguity aversion. At face value, this finding seems to contradict our result on the corresponding correlation being positive. However, this is not the case. Making the information that participants receive less reliable, a variation of our original experimental setup is able to replicate the finding of Shishkin and Ortoleva (2023). This establishes that the reliability of the information is key for the result on the negative value of information.

A large strand of the literature documents the prevalence of real-world instances of information aversion (Hertwig and Engel, 2016; Golman, Hagmann, and Loewenstein, 2017; Brown and Walasek, 2020), running from behavioral finance (Karlsson, Loewenstein, and Seppi, 2009) to health (Oster, Shoulson, and Dorsey, 2013; Ho, Hagmann, and Loewenstein, 2021) and managerial decision-making (Deshpande and Kohli, 1989). Other experimental studies checked whether subjects are information averse in games, studying the phenomenon referred to as strategic ignorance (Dana, Weber, and Kuang, 2007; Grossman and van der Weele, 2017). Finally, a number of studies proposed different reasons for information aversion such as anxiety (Kőszegi, 2003; Eliaz and Spiegler, 2006; Epstein, 2008), disappointment aversion (Dillenberger, 2010; Andries and Haddad, 2020), regret aversion (Krähmer and Stone, 2013; Somasundaram and Diecidue, 2017), optimism maintenance (Brunnermeier and Parker, 2005; Oster, Shoulson, and Dorsey, 2013) or belief investments (Jonas et al., 2001; Dana, Weber, and Kuang, 2007). Our study in this paper implemented as a controlled and abstract laboratory experiment highlights the robustness of the phenomenon of information aversion and suggests in ambiguity aversion a potential driver for such information averse behavior.

Section 2 provides formal definitions of the tested concepts. The next two sections present our experimental designs and results. Section 5 provides a discussion of the underlying assumptions and design choice. The final section concludes.

2 Value of Information, Ambiguity Attitude, and Dynamic Consistency

Let S be a set of states. For simplicity, assume that S is finite. Subsets of S are referred to as events, i.e., any $E \subseteq S$ is an event. Let $X \subseteq \mathbb{R}$ denote a set of outcomes. An act f is a function from S to X. We consider a DM who has preferences $\succcurlyeq$ over the set $\mathcal{F}$ of all possible acts. The DM's preferences conditional on the occurrence of some event $E \subset S$ are denoted by $\succcurlyeq_E$.

From the ex-ante view, information about the state of the world can be represented as a partition of S (Aumann, 1974; Gilboa and Lehrer, 1991). For example, note that the information on the color of the randomly drawn ball in the setting described in the Introduction may be represented as the partition of $\{1,2,3,4\}$ into the two events $G = \{1,3\}$ and $B = \{2,4\}$. In what follows, we compare the theoretical value of information in two models of decision-making, one without ambiguity, the other with ambiguity.

2.1 Value of Information in SEU Model

In the SEU model (Savage, 1954), the DM's preferences are characterized by a utility function $u : X \to \mathbb{R}$ and a unique subjective belief p about the states of the world in S. Assume that the DM is facing the problem of choosing a single act from a finite set $F \subseteq \mathcal{F}$. The value V_0 of such a choice set without information is given by

$$V_0 = \max_{f \in F} \sum_{s \in S} p(s) u\left(f(s)\right) \tag{1}$$

i.e., the maximal expected utility available by choosing an act from F. The value V_E of the same choice set conditional on an event $E \subset S$ is given by

$$V_E = \max_{f \in F} \sum_{s \in E} p_E(s) u\left(f(s)\right) \tag{2}$$

which is the maximal expected utility with respect to the update p_E of p. An elementary piece of information is represented by a partition of S into two sets $A, \bar{A} \subset S$. Then, the value of information is defined in the following way.

Definition 2.1 (Blackwell (1953))**.** *Let $F \subseteq \mathcal{F}$ be a choice set and $\{A, \bar{A}\}$ be a partition of S. Under SEU, the value of information $\{A, \bar{A}\}$ is given by*

$$VOI = \left[p(A)V_A + p(\bar{A})V_{\bar{A}}\right] - V_0 \tag{3}$$

where V_0 and $V_A, V_{\bar{A}}$ are defined by equations (1) and (2) respectively.

Thus, the value of information equals the expected increase in utility resulting from being able to condition the choice of an act on the elements of the partition. Note that because $p = p(A)p_A + p(\bar{A})p_{\bar{A}}$ and mathematical expectation is linear in probabilities, we have

$$p(A) \sum_{s \in A} p_A(s) u\left(f(s)\right) + p(\bar{A}) \sum_{s \in \bar{A}} p_{\bar{A}}(s) u\left(f(s)\right) = \sum_{s \in S} p(s) u\left(f(s)\right)$$

for any $f \in F$. Therefore, the expression in the square brackets of (3) is always greater than or equal to V_0, because the optimal acts associated with V_A and $V_{\bar{A}}$ are chosen independently.

Thus, $VOI \geq 0$ for any DM, choice set, and information under SEU. Intuitively, information expands the DM's choice set to include combinations of conditional acts, which cannot harm.

2.2 Value of Information in Maxmin Expected Utility Model

Under ambiguity, however, the value of information can be negative. According to the Maxmin Expected Utility (MEU) model (Gilboa and Schmeidler, 1989; Ivanenko and Labkovskii, 1986), the DM's subjective beliefs form a set $\mathcal{C}$ of probability distributions on the state space. The DM ranks acts by maximizing the minimal expected utility of an act with respect to this set of beliefs. Hence, the value of a choice set $F \subseteq \mathcal{F}$ without information is

$$V_0 = \max_{f \in F} \min_{p \in \mathcal{C}} \sum_{s \in S} p(s) u\left(f(s)\right) \tag{4}$$

and the value of the same choice set conditional on an event $E \subset S$ is

$$V_E = \max_{f \in F} \min_{p_E \in \mathcal{C}_E} \sum_{s \in E} p_E(s) u\left(f(s)\right) \tag{5}$$

where $\mathcal{C}_E$ is the Full Bayesian (includes updates of all beliefs) update of $\mathcal{C}$.

Definition 2.2 (Li (2020)). *Let $F \subseteq \mathcal{F}$ be a choice set and $\{A, \bar{A}\}$ be a partition of S. Under MEU, the value of information $\{A, \bar{A}\}$ is given by*

$$VOI = \min_{p \in \mathcal{C}} \left[p(A)V_A + p(\bar{A})V_{\bar{A}}\right] - V_0 \tag{6}$$

where V_0 and $V_A, V_{\bar{A}}$ are defined by equations (4) and (5) respectively.

To see that the value of information in (6) can be negative, the following intuition may be helpful. Although the DM can choose different acts conditional on A and $\bar{A}$, the minima over $\mathcal{C}_A$ and $\mathcal{C}_{\bar{A}}$ in (5) may be reached at the updates of different beliefs. Moreover, there is another minimum in (6). Hence, information works as if it expands both the DM's and Nature's choice sets. Whether the value of information is negative, positive, or zero depends on which side benefits from such an expansion more.

From the experimental perspective, estimating the values in (6) can be difficult, as it requires eliciting the DM's set of beliefs. However, as will be explained in the next section, the structure of our experiment allows us to assume that $V_A = V_{\bar{A}}$. Hence, the expression in (6) becomes

$$VOI = V_E - V_0 \tag{7}$$

for $E \in \{A, \bar{A}\}$. Thus, we get a simple criterion for the value of information to be negative, which is $VOI < 0$ if and only if $V_E < V_0$.

2.3 Ambiguity Attitude and Dynamic Consistency

Ellsberg (1961) pointed out that missing information about probabilities, commonly referred to as ambiguity, affects preferences in a way that cannot be explained by the SEU model. Consider three events represented by non-empty and pairwise disjoint sets $A, B, C \subset S$. Assume that only the probabilities of A and $B \cup C$ are known. Let b_E be a bet on an event $E \subseteq S$, i.e., $b_E(s) = x$ for $s \in E$ and $b_E(s) = y$ otherwise, for some $x \succ y$. Then, preferences $b_A \succ b_B$ and $b_{A\cup C} \prec b_{B\cup C}$ are inconsistent with SEU and are referred to as ambiguity averse. The opposite preferences $b_A \prec b_B$ and $b_{A\cup C} \succ b_{B\cup C}$, referred to as ambiguity seeking, cannot be described by SEU either. The remaining two variants $b_A \succ b_B$, $b_{A\cup C} \succ b_{B\cup C}$ and $b_A \prec b_B$, $b_{A\cup C} \prec b_{B\cup C}$ are consistent with SEU and referred to as ambiguity neutral. For a summary of experimental work on this topic, see Oechssler and Roomets (2015) and Trautmann and Van De Kuilen (2015).

Dynamic consistency links conditional and unconditional preferences in the following way.

Definition 2.3. *Let $\{A, \bar{A}\}$ be a partition of S. Preferences are dynamically consistent if for any $f, g \in \mathcal{F}$, we have $f \succcurlyeq g$ only if either $f \succcurlyeq_A g$ or $f \succcurlyeq_{\bar{A}} g$.*

For example, the preferences $(f \succcurlyeq g,\, g \succ_A f,\, g \succ_{\bar{A}} f)$ are dynamically inconsistent because the ex-ante preference of $f \succcurlyeq g$ is not implemented in any of the events in $\{A, \bar{A}\}$, since both $g \succ_A f$ and $g \succ_{\bar{A}} f$. While dynamic consistency is implied by the combination of SEU and the standard updating rules, this is not a given under ambiguity models such as MEU.

3 Experimental Design

We conducted two experiments with an overall of 222 incentivized participants in the laboratory. The first experiment provides a simple test of the negative value of information and its relation to ambiguity aversion. Complementing this experiment, we ran a second experiment to probe the robustness of the results from the first experiment. This second experiment extends our analysis of the negative value of information to its sensitivity on the stakes in the experiment, to information with unknown or ambiguous reliability and to the role of complexity in explaining the results.

3.1 Design of First Experiment

Our first experiment builds on the variant of an Ellsberg urn depicted in Figure 2. The urn contains 21 balls. Every ball has a color and is marked with a number. There are five green balls marked with the number 1 and five blue balls marked with the number 2, $n_1 = n_2 = 5$.

Each of the remaining eleven balls is either green and marked with a 3, or, it is blue and marked with a 4. The exact numbers n_3 and n_4 are unknown to the participants. But, it is known that $n_3 + n_4 = 11$. Subjects in our experiment are asked to choose between different bets on the number a randomly drawn ball from the urn is marked with.

Figure 2: Ellsberg-type urn

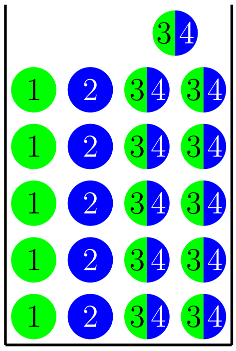

Our experiment consists of five problems of pairwise choice between bets. Out of these five choice problems, three were designed to test information aversion and dynamic consistency under uncertainty. The other two choice problems elicit participants' ambiguity preferences. In each choice problem, participants are asked to choose between two different bets. Each choice problem reveals how much each bet pays depending on the number of a randomly drawn ball from the urn.

3.1.1 Testing Value of Information

Intuitively, information that breaks an existing hedge against ambiguity may lead to a negative value of information for ambiguity averse DMs.

To see how we translated this intuition into our experimental design, consider the choice between the two bets on a randomly drawn ball from the urn, listed in Table 1. When the color of the drawn ball is not known upfront, each of these bets is risky. The probability of getting the higher prize is $\frac{10}{21}$ for $\overline{\mathbf{12}}$ and $\frac{11}{21}$ for $\overline{\mathbf{34}}$. On the other hand, when the color is known upfront, each of these bets becomes ambiguous. For example, after knowing that the drawn ball is green, the winning chance in $\overline{\mathbf{12}}$ changes from $\frac{10}{21}$ to $\frac{5}{5+n_3}$, where n_3 is unknown. Similarly, the winning chance in $\overline{\mathbf{34}}$ changes from $\frac{11}{21}$ to $\frac{n_3}{5+n_3}$, which is ambiguous as well. For an ambiguity averse DM such ambiguity is undesirable and he may even be willing to forego a small payoff premium to avoid the ambiguity that comes with the information about the color. A similar argument can be made in case the drawn ball is blue.

Next, we lay out the precise formulation of our decision tasks eliciting the value of information. In Problem 1, participants are asked to choose between the two bets $\overline{\mathbf{12}}$ and $\overline{\mathbf{34}}$ not knowing

Table 1: Bets in Problem 1 (No Color Info condition)

	5 balls	5 balls	11 balls	
	1	2	3	4
$\overline{\mathbf{12}}$	€10	€10	€4	€4
$\overline{\mathbf{34}}$	€4	€4	€10	€10

the color of the drawn ball (No Color Info condition). Table 1 specifies what each bet pays depending on the number that the randomly drawn ball is marked with.

For Problem 2, a ball is randomly drawn before participants are asked to make their decision between the two bets of this choice problem. Participants can see the color of the drawn ball, but the number it is marked with remains unknown to them (Color Info condition). Table 2 specifies what each bet pays depending on the number. Note that the only difference between the bets in this and the previous problem lies in the payoff increase of €0.5 in every state.

Table 2: Bets in Problem 2 (Color Info condition)

	5 balls	5 balls	11 balls	
	1	2	3	4
$\overline{\mathbf{12}}^*$	€10.5	€10.5	€4.5	€4.5
$\overline{\mathbf{34}}^*$	€4.5	€4.5	€10.5	€10.5

In Problem 3, participants are asked to choose between the two bets they have chosen in the previous two problems. Let $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ and $c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*)$ specify participant's bets chosen in Problems 1 and 2. Then, in the third problem, this participant is asked to choose between

$$c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}}) \qquad \text{vs.} \qquad c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*)$$

Say, a participant has chosen $\overline{\mathbf{34}}$ in Problem 1, i.e., $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}}) = \overline{\mathbf{34}}$, and $\overline{\mathbf{12}}^*$ in Problem 2, i.e., $c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*) = \overline{\mathbf{12}}^*$. Then, in the third problem, she is asked to choose between $\overline{\mathbf{34}}$ and $\overline{\mathbf{12}}^*$. Note that $\overline{\mathbf{34}}$ still refers to the No Color Info condition and $\overline{\mathbf{12}}^*$ still refers to the Color Info condition.

To apply the framework of Section 2, we make the following assumption.

Assumption 1. (MEU) *The DM's preferences are represented by the MEU model with the state space $S = \{1, 2, 3, 4\}$, the set of outcomes $X = \{4, 4.5, 10, 10.5\}$ and acts $\overline{\mathbf{12}}$, $\overline{\mathbf{34}}$, $\overline{\mathbf{12}}^*$, $\overline{\mathbf{34}}^*$. The DM's utility function $u : X \to \mathbb{R}$ is strictly increasing. The information is represented by the partition $\{G, B\}$, where $G = \{1, 3\}$ and $B = \{2, 4\}$.*

We classify participants choosing bet $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ in Problem 3 as showing a negative value of information and those choosing $c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*)$ as showing a non-negative value of information.

This classification is based on the following reasoning. Note that, since act $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ is optimal in the set $F = \{\overline{\mathbf{12}}, \overline{\mathbf{34}}\}$ without information, the value of $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ is equal to V_0 in Equation (4). Similarly, the value of $c_2(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ (not observed in the experiment) equals V_E from Equation (5) for $E = G$ or $E = B$ depending on what ball was actually drawn. Since acts $\overline{\mathbf{12}}$ and $\overline{\mathbf{34}}$ are symmetric with respect to color, as is the information about urn compositions, this motivates the following assumption.

Assumption 2. (Color Symmetry) $V_G = V_B$.

In other words, the DM's choice in Problem 2 and the resulting utility do not depend on the color of the drawn ball. Assumption 2 allows us to simplify the value of information given by Equation (6) to $VOI = V_E - V_0$, where $E \in \{G, B\}$ is the color of the drawn ball. Because acts $\overline{\mathbf{12}}^*$ and $\overline{\mathbf{34}}^*$ dominate acts $\overline{\mathbf{12}}$ and $\overline{\mathbf{34}}$ by €0.5 in every state, we have $c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*) \succ c_2(\overline{\mathbf{12}}, \overline{\mathbf{34}})$. Hence, the observed preference $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}}) \succ c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*)$ implies $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}}) \succ c_2(\overline{\mathbf{12}}, \overline{\mathbf{34}})$. The latter implies $V_E < V_0$, from which $VOI < 0$ follows.

Note that if the observed preference $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}}) \succ c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*)$ is in fact the indifference $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}}) \sim c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*)$, the value of information is still strictly negative because we have $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}}) \sim c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*) \succ c_2(\overline{\mathbf{12}}, \overline{\mathbf{34}})$, which implies $V_E < V_0$.

Intuitively, $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}}) \succ c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*)$ means that the DM leaves money on the table in exchange for not knowing the color of the randomly drawn ball upfront. Since all bets are symmetric w.r.t. color, this observation does not depend on whether the color of the ball which was actually drawn is green or blue. Section 5.1 provides a critical discussion of the two assumptions laid out here.

3.1.2 Testing Dynamic Consistency

Dynamic consistency requires that ex-ante contingent choices are respected by updated preferences. In our experiment, we classify participants choosing bet $\overline{\mathbf{12}}$ from Problem 1 and bet $\overline{\mathbf{12}}^*$ from Problem 2, or respectively, bets $\overline{\mathbf{34}}$ and $\overline{\mathbf{34}}^*$ as dynamically consistent. All other participants, i.e., those choosing $\overline{\mathbf{12}}$ and $\overline{\mathbf{34}}^*$, or, $\overline{\mathbf{34}}$ and $\overline{\mathbf{12}}^*$, we classify as dynamically inconsistent. The assumption underlying this classification is the following.

Assumption 3. (Color Symmetry Continued) *The DM's conditional preference $\succsim_E$ between acts $\overline{\mathbf{12}}^*$ and $\overline{\mathbf{34}}^*$ does not depend on the event $E \in \{G, B\}$.*

Say, a green ball is drawn in the Color Info condition. To see why preferences

$$\overline{\mathbf{12}} \succ \overline{\mathbf{34}},\ \overline{\mathbf{12}}^* \prec_G \overline{\mathbf{34}}^*$$

are dynamically inconsistent, note that, abusing notation, $\overline{\mathbf{12}}^* = \overline{\mathbf{12}} + 0.5$ and $\overline{\mathbf{34}}^* = \overline{\mathbf{34}} + 0.5$. It follows from Assumption 1 that $\overline{\mathbf{12}} \succ \overline{\mathbf{34}}$ implies $\overline{\mathbf{12}}^* \succ \overline{\mathbf{34}}^*$. On the other hand, $\overline{\mathbf{12}}^* \prec_G \overline{\mathbf{34}}^*$ implies $\overline{\mathbf{12}}^* \prec_B \overline{\mathbf{34}}^*$ by Assumption 3. Taken together, the preferences above imply

$$\overline{\mathbf{12}}^* \succ \overline{\mathbf{34}}^*,\ \overline{\mathbf{12}}^* \prec_G \overline{\mathbf{34}}^*,\ \overline{\mathbf{12}}^* \prec_B \overline{\mathbf{34}}^*,$$

which are dynamically inconsistent by Definition 2.3. The same reasoning also establishes that the preferences for $\overline{\mathbf{34}}$ and $\overline{\mathbf{12}}^*$ are dynamically inconsistent, as well.

3.1.3 Testing Ambiguity Attitude

The structure of our test for ambiguity aversion is a variation of the Ellsberg 3-color problem. Following the ambiguity-literature, we test for the DM's ambiguity attitude by offering two choice problems. Specifically, in Problems 4 and 5, participants choose between two bets in the No Color Info condition (see Table 3). Note that the bets $\overline{\mathbf{1}}$ and $\overline{\mathbf{14}}$, as well as $\overline{\mathbf{3}}$ and $\overline{\mathbf{34}}$, differ only in their payoffs in state 4. While $\overline{\mathbf{1}}$ and $\overline{\mathbf{34}}$ are risky bets, the bets $\overline{\mathbf{3}}$ and $\overline{\mathbf{14}}$ are both ambiguous. By the usual interpretation, participants choosing the risky bets $\overline{\mathbf{1}}$ and $\overline{\mathbf{34}}$ (resp., the ambiguous bets $\overline{\mathbf{3}}$ and $\overline{\mathbf{14}}$) are classified as ambiguity averse (resp., seeking). All other participants, i.e., those choosing $\overline{\mathbf{1}}$ and $\overline{\mathbf{14}}$, or, $\overline{\mathbf{3}}$ and $\overline{\mathbf{34}}$, are classified as ambiguity neutral.

Table 3: Bets in Problems 4 and 5 (No Color Info condition)

	5 balls	5 balls	11 balls	
	1	2	3	4
$\overline{\mathbf{1}}$	€10	€4	€4	€4
$\overline{\mathbf{3}}$	€4	€4	€10	€4
$\overline{\mathbf{14}}$	€10	€4	€4	€10
$\overline{\mathbf{34}}$	€4	€4	€10	€10

A problem common to ambiguity related experiments is how to deal with indifference. We tackle this problem in two ways. First, the number $n_3 + n_4$ of ambiguous balls is one larger than the number $n_1 + n_2$ of risky balls. This extra ball acts as a tie breaker. Hence, preference for $\overline{\mathbf{1}}$ over $\overline{\mathbf{3}}$ cannot be explained by indifference alone. Second, for each choice problem, we asked participants about their confidence in their choices such that we could interpret participants stating the lowest level of confidence as having no confidence that their choice is better than the alternative in this choice problem, i.e., as being indifferent between the two options. In Section 4.2.7, we use this measure as a robustness check in the analysis of our results.

Finally, Table 4 provides a summary chart of all five experimental tasks that spells out how choices in each task constitute the tests of the value of information, dynamic consistency, and ambiguity attitude in our experiment. Appendix A contains the analysis of our five choice problems using different models of decision-making.

Table 4: Summary chart of experimental problems and tests

	Problem 1	Problem 2	Problem 3	Problem 4	Problem 5
	$\overline{\mathbf{12}}$ vs. $\overline{\mathbf{34}}$	$\overline{\mathbf{12}}^*$ vs. $\overline{\mathbf{34}}^*$	$c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ vs. $c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*)$	$\overline{\mathbf{1}}$ vs. $\overline{\mathbf{3}}$	$\overline{\mathbf{14}}$ vs. $\overline{\mathbf{34}}$
Test					
Negative VOI	-	-	$c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}})$	-	-
Dynamically Consistent	$\overline{\mathbf{12}}$ $\overline{\mathbf{34}}$	$\overline{\mathbf{12}}^*$ $\overline{\mathbf{34}}^*$	- -	- -	- -
Dynamically Inconsistent	$\overline{\mathbf{12}}$ $\overline{\mathbf{34}}$	$\overline{\mathbf{34}}^*$ $\overline{\mathbf{12}}^*$	- -	- -	- -
Ambiguity Averse	-	-	-	$\overline{\mathbf{1}}$	$\overline{\mathbf{34}}$
Ambiguity Neutral	- -	- -	- -	$\overline{\mathbf{1}}$ $\overline{\mathbf{3}}$	$\overline{\mathbf{14}}$ $\overline{\mathbf{34}}$
Ambiguity Seeking	-	-	-	$\overline{\mathbf{3}}$	$\overline{\mathbf{14}}$

Notes: Problems 1, 4, and 5 refer to the No Color Info condition, Problem 2 refers to the Color Info condition, Problem 3 refers to both.

3.1.4 Implementation & Lab Procedures of First Experiment

The experiment was conducted in November and December 2019 in the AWI Lab at Heidelberg University. We implemented the above-described decision tasks as a pen-and-paper experiment. Subjects were recruited via SONA System and paid in cash directly after the experiment. All participants received a show-up fee of €4 and could earn additionally up to €10.5 from the decision tasks. The experiment took about 45 minutes, for which participants earned, on average, €12.8. For practical reasons, in the laboratory, we used non-transparent, colored balls that could be opened and filled each with a folded piece of paper that was marked with a number from one to four. Furthermore, instead of urns we used cardboxes. We used two identical boxes, i.e., one box for the No Color Info condition and another for the Color Info condition. Before each session, the boxes were checked to contain the correct distribution of colored, marked balls. Participants did not have any information as to what the distribution of the eleven ambiguous balls in each box were, but the physical boxes were visibly placed on the experimenter's table for all subjects to see and could be inspected by participants after each session. Before decision sheets were distributed, uncertainty about the ball color of the randomly drawn ball for the Color Info condition was resolved. This was done with the help of a randomly selected participant who performed this random draw. The ball's color was announced and the ball remained unopened on top of the cardbox, for everyone to see, until the end of the experiment. So, the number it was marked with remained unknown. Participants marked their choices on the decision sheets and answered a demographic questionnaire including our question about urn-compositions. Then the decision sheets were collected. Another randomly selected participant randomly picked one of five cards that determined which of the five choice problems was payoff-relevant. The shuffled cards were presented face down. For each of the five choice problems, there was one card referring to it. If the payoff-relevant choice problem involved the Color Info condition (Problems 2 and 3), the previously drawn ball was opened and the number was revealed to participants. If the payoff-relevant choice problem involved the

No Color Info condition (Problems 1, 3, 4, and 5), the same participant performed a random draw from the other box, opened the drawn ball, and revealed the number it was marked with to participants. Final payoffs were calculated, participants were paid and dismissed from the lab. On their way out, they could inspect the urns if they wanted to.[2]

3.2 Design of Second Experiment

The second experiment complements the first one. Their designs overlap in many respects. But, they also differ in subtle details. The similarities and differences in their designs can be understood best by dividing the questions of the second experiment into the following three categories: replication (REP) questions, complexity (COM) questions and ambiguous reliability (ARL) questions. Overall, the second experiment comprises ten choice problems. Table 5 shows how the ten choice problems of the second experiment divide into the three categories of REP, COM, and ARL questions. Each of the ten choice problems is implemented in a Multiple Price List (MPL) format. This amounts to fifteen questions (or lines) per every choice problem. The middle (8th) question of every choice problem always offers a choice between the two plain urn-bets of the corresponding problem. The first seven questions each add a different extra payment to the first plain urn-bet, but keep the second plain urn-bet as is. The last seven questions do it vice versa, by keeping the first plain urn-bet as is and adding extra payments to the second plain urn-bet. These added extra payments take the values of €5.00, €2.00, €1.50, €1.00, €0.75, €0.50, and €0.25.

Table 5: Summary chart of the second experiment and its robustness tests

	Problem 1	Problem 2	Problem 3	Problem 4	Problem 5	Problem 6	Problem 7	Problem 8	Problem 9	Problem 10
Questions										
REP	X	X	X	X	X	-	-	-	-	-
COM	-	-	-	-	-	X	X	X	-	-
ARL	-	-	(X)	-	-	-	-	-	X	X

Notes: All ten choice problems of the second experiment are implemented using an MPL format. Problems 1, 2, 3, 4, and 5 replicate (REP) the five choice problems from the first experiment. Problems 6, 7, and 8 provide a test for complexity (COM) avoidance in the test of the negative value of information. Problems 9 and 10 vis-a-vis Problem 3 test for the role of the reliability of the information in the test of the negative value of information.

3.2.1 Replication (REP) Questions

The questions in this category aim at a replication of the first experiment. They involve the same tests as the first experiment does. Also, the two urns for these tests, Urn U_1 and Urn K_1, are as described in Section 3.1. Where the replication questions differ from those of the first experiment is in the way they are implemented. Instead of asking a single binary question per

[2] A replication package including all instructions, data, analysis files and other materials is available at https://osf.io/w8k2p/. Instructions were translated from German. Original instructions are available from the authors upon request.

every choice problem, each of the five choice problems is now presented in the MPL format laid out above. Table 6 lists the plain urn-bets of Problems 1, 2, 3, and 4.

Table 6: Plain urn-bets without extra payments (i.e., middle questions) of Problems 1, 2, 3, and 4 in the second experiment

	5 balls	5 balls	11 balls	
	1	2	3	4
$\overline{\mathbf{1}}$	€10	€4	€4	€4
$\overline{\mathbf{3}}$	€4	€4	€10	€4
$\overline{\mathbf{14}}$	€10	€4	€4	€10
$\overline{\mathbf{34}}$	€4	€4	€10	€10
$\overline{\mathbf{12}}$	€10	€10	€4	€4
$\overline{\mathbf{34}}$	€4	€4	€10	€10
$\overline{\mathbf{12}}^*$	€10	€10	€4	€4
$\overline{\mathbf{34}}^*$	€4	€4	€10	€10

Notes: No Color info condition (all Urn U_1): Problem 1 ($\overline{\mathbf{1}}$ vs. $\overline{\mathbf{3}}$), Problem 2 ($\overline{\mathbf{14}}$ vs. $\overline{\mathbf{34}}$), Problem 3 ($\overline{\mathbf{12}}$ vs. $\overline{\mathbf{34}}$); Color Info condition (Urn K_1): Problem 4 ($\overline{\mathbf{12}}^*$ vs. $\overline{\mathbf{34}}^*$)

The plain urn-bets of Problem 5 are determined by the participant's choice in the middle question of Problem 3 and the choice in the middle question of Problem 4. The bottom rows of Table 6 list the corresponding possible plain urn-bets. Let $c_3(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ and $c_4(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*)$ specify a participant's bets chosen in these middle questions of Problems 3 and 4. Then, $c_3(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ and $c_4(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*)$ serve as the plain urn-bets for the implementation of Problem 5 in the MPL format.

3.2.2 Complexity (COM) Questions

The questions in this category serve as a robustness check for the results on the negative value of information from the first experiment and the REP questions of the second experiment. The idea is to test whether these results could be driven by the relative complexity of the choice problems where the color of the drawn ball is known.

To check this, the setup of the complexity questions is almost identical to that of the replication questions on the negative value of information. The only difference between the two setups lies in the different Ellsberg-type urns to which the questions refer. Recall that for the replication questions, the Ellsberg-type urns are as described in Section 3.1. The number of 3 balls in these urns is unknown, as is the way in which this number was determined. There is ambiguity. This is different for the two urns, Urn U_2 and Urn K_2, in the complexity category. For these urns, it is known how this number is determined such that there is no ambiguity. To be precise, the number of 3 balls (and, thus, also the number of 4 balls, $n_4 = 11 - n_3$) in the Ellsberg-type urns for the complexity questions is determined randomly such that every number 0, 1, 2, ..., 11 is equally likely. Everything else about these urns is the same as for the urns in the replication category, i.e., as described in Section 3.1. This way, the complexity questions retain the complexity of the replication questions, but not the ambiguity about how the composition of the urn was determined. Of course, the instructions explicitly inform participants about the differences between the two categories. Table 7 lists the plain urn-bets of Problems 6 and 7.

Table 7: Plain urn-bets without extra payments (i.e., middle questions) of Problems 6 and 7 in the second experiment

	5 balls	5 balls	11 balls	
	1	2	3	4
$\overline{\mathbf{12}}^{\dagger}$	€10	€10	€4	€4
$\overline{\mathbf{34}}^{\dagger}$	€4	€4	€10	€10
$\overline{\mathbf{12}}^{\dagger\dagger}$	€10	€10	€4	€4
$\overline{\mathbf{34}}^{\dagger\dagger}$	€4	€4	€10	€10

Notes: No Color info condition (Urn U_2): Problem 6 ($\overline{\mathbf{12}}^{\dagger}$ vs. $\overline{\mathbf{34}}^{\dagger}$); Color Info condition (Urn K_2): Problem 7 ($\overline{\mathbf{12}}^{\dagger\dagger}$ vs. $\overline{\mathbf{34}}^{\dagger\dagger}$)

The plain urn-bets of Problem 8 are determined by the participant's choice in the middle question of Problem 6 and in the middle question of Problem 7. Table 7 lists all possible plain urn-bets. Let $c_6(\overline{\mathbf{12}}^{\dagger}, \overline{\mathbf{34}}^{\dagger})$ and $c_7(\overline{\mathbf{12}}^{\dagger\dagger}, \overline{\mathbf{34}}^{\dagger\dagger})$ specify a participant's bets chosen in these middle questions of Problems 6 and 7. Then, $c_6(\overline{\mathbf{12}}^{\dagger}, \overline{\mathbf{34}}^{\dagger})$ and $c_7(\overline{\mathbf{12}}^{\dagger\dagger}, \overline{\mathbf{34}}^{\dagger\dagger})$ serve as the plain urn-bets for the implementation of Problem 8 in the MPL format.

3.2.3 Ambiguous Reliability (ARL) Questions

The questions in this category extend our analysis of the negative value of information to information with unknown or ambiguous reliability.

To do so, we set up an urn which is identical to the corresponding one in the replication category. Figure 2 serves again as illustration. Then, we put each of the 21 balls in the urn inside another slightly larger ball. Figure 3 illustrates the resulting Urn U_3. Each of the larger balls is non-transparent, openable, and also of green or blue color. The key idea is that the color of inner and outer ball do not have to match. When a ball from this urn is drawn, participants only receive information about the color of the outer ball. The color of the inner ball remains unknown. Therefore, knowing the color of the outer ball, say green, does not rule out that the inner ball is blue, containing an even number (2 or 4). It is in this way that the information here is made less reliable. The reason being that, as for the other categories, all questions in the ARL category refer to the number that a drawn ball is marked with and this number only has a guaranteed relationship to the color of the inner ball, not to that of the outer ball.

In essence, the setup here involves a variation of Problem 4 and Problem 5 from the REP category of the second experiment. Table 8 lists the plain urn-bets of Problem 9.

Figure 3: Ellsberg-type Urn U_3 with balls inside other balls

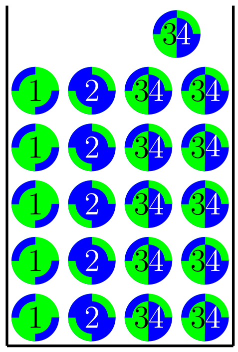

Table 8: Plain urn-bets without extra payments (i.e., middle questions) of Problem 9 in the second experiment

	5 balls	5 balls	11 balls	
	1	2	3	4
$\overline{\mathbf{12}}^{\circ}$	€10	€10	€4	€4
$\overline{\mathbf{34}}^{\circ}$	€4	€4	€10	€10

Notes: Outer Color Info condition (Urn U_3): Problem 9 ($\overline{\mathbf{12}}^{\circ}$ vs. $\overline{\mathbf{34}}^{\circ}$)

The plain urn-bets of Problem 10 are determined by the participant's choice in the middle question of Problem 3 and the choice in the middle question of Problem 9. Table 6 and Table 8 list the corresponding possible plain urn-bets. Let $c_3(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ and $c_9(\overline{\mathbf{12}}^{\circ}, \overline{\mathbf{34}}^{\circ})$ specify a participant's bets chosen in these middle questions of Problems 3 and 9. Then, $c_3(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ and $c_9(\overline{\mathbf{12}}^{\circ}, \overline{\mathbf{34}}^{\circ})$ serve as the plain urn-bets of Problem 10.

3.2.4 Implementation & Lab Procedures of Second Experiment

The second experiment was conducted in March 2023 in the BEElab of the School of Business and Economics at Maastricht University. To keep it as close as possible to the first experiment, we again implemented the above-described decision tasks as a pen-and-paper experiment. Subjects were recruited via ORSEE System and paid via bank transfer after the experiment. All participants received a show-up fee of €3 and could earn additionally up to €15 from the decision tasks. The experiment took about 50 minutes, for which participants earned, on average, €11.9. As in the first experiment, the equipment consisted of non-transparent, openable, colored balls and cardboxes. Before each session, the overall five identical cardboxes were checked to contain the correct distribution of colored, marked balls. The two boxes for the condition replicating the first experiment were filled as in the first experiment (see Section 3.1.4). For the two cardboxes of the complexity condition, the number of 3 balls was determined using a random number generator. The two boxes were then filled accordingly with the corresponding number of green balls containing folded pieces of paper, each marked with the number 1 or 3, and the corresponding number of blue balls, each marked in the same way with the number 2 or 4. For the condition on information with unknown reliability, we prepared the balls as in the condition replicating the first experiment. These balls were then put inside slightly larger, non-transparent and openable balls. The larger balls were of green and blue color, as well. True to the nature of this condition, participants did not have any information as to how the color of the larger ball relates to the color of the ball inside it. The physical boxes were placed in the center of the room such that every subject could see them throughout the entire experiment. Participants were allowed to inspect the cardboxes after the experiment. Before decision sheets were distributed, it was emphasized that only one question will determine earnings. Then, uncertainty about the ball color of the randomly drawn ball from the corresponding boxes in every Color Info condition was resolved. This was done with the help of a randomly selected participant who performed these random draws. The ball's color was announced and it remained unopened on top of the cardbox, for everyone to see, until the end of the experiment. So, the number it was marked with remained unknown. Participants marked their choices on the decision sheets and answered a demographic questionnaire including our question about urn-compositions. The order of decision sheets and questionnaire was reversed for half of the participants. Then all sheets were collected and uncertainty about which of the questions determined participants' payoffs was resolved. This was done by use of a mobile app that generates random numbers and the help of another randomly selected participant who operated the app. If the payoff-relevant question involved a box with a previously drawn ball, this ball was opened and the number was revealed to participants. If the payoff-relevant question involved another box, then the same participant performed a random draw from the corresponding box, opened the drawn ball, and revealed the number it was marked with to participants. Final payoffs were calculated and participants were dismissed from the lab.[3]

4 Results

4.1 Results of First Experiment

4.1.1 Summary Statistics

In total, 115 subjects participated in the first experiment. Subjects' average age was 23.2 years, and the share of economics students was 25%. With 57%, the share of female participants was reasonably close to 50% in our experiment.

[3]A replication package including all instructions, data, analysis files and other materials is available at https://osf.io/w8k2p/.

Figure 4: %-Choices for the two Options in Problems 1-5

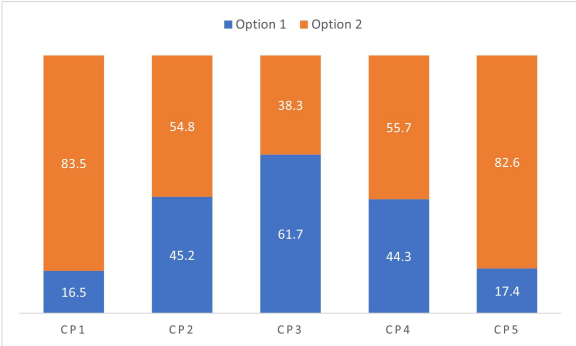

Figure 4 shows what fraction of participants chose which of the two options in each of our five choice problems. For example, 17% chose the stochastically dominated Option 1 ($\overline{\mathbf{12}}$) in Problem 1 (see Section 4.2.7 for a robustness check that excludes these subjects).

4.1.2 Negative Value of Information

Our main research question is whether subjects do assign a negative value to information under uncertainty. Such participants would choose their preferred option in Problem 1 over their preferred option in Problem 2, when it comes to their choice between these two bets in Problem 3. Table 9 shows the percentage of participants that made negative VOI-choices. The main result of our study is very clear: **at 62%, the share of participants showing a negative value of information is substantial.**[4] We refer to these 71 subjects as the VOI− group and to the remaining 44 subjects as the VOI+ group.

Table 9: Share of the negative VOI

Problem 3	Share	N
$c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ vs. $c_2(\overline{\mathbf{12}}^*, \overline{\mathbf{34}}^*)$	62% VOI− **	115

Notes: Column "Share" shows percentage of $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ chosen in Problem 3. Binary choice: Two-sided binomial test against $p = 0.5$. ** denotes significance at 5%.

Motivated by recent theoretical work, we checked whether the negative value of information is correlated with ambiguity aversion. Indeed, as Figure 5a shows, **subjects in the VOI− group were more often ambiguity averse than subjects in the VOI+ group** ($p = 0.028$, two-sided proportion test). The percentage doubles, from 21% to 42%.[5] On the other hand, the VOI+ group was more often ambiguity neutral than the VOI− group ($p = 0.080$, two-sided proportion test).

Among ambiguity averse subjects, 77% showed a negative value of information (Figure 5b).

[4]This share is 64% if subjects who chose the stochastically dominated option in Problem 1 are excluded.

[5]These shares are 20% and 44% if subjects violating first-order stochastic dominance are excluded ($p = 0.03$, two-sided proportion test, $N = 96$).

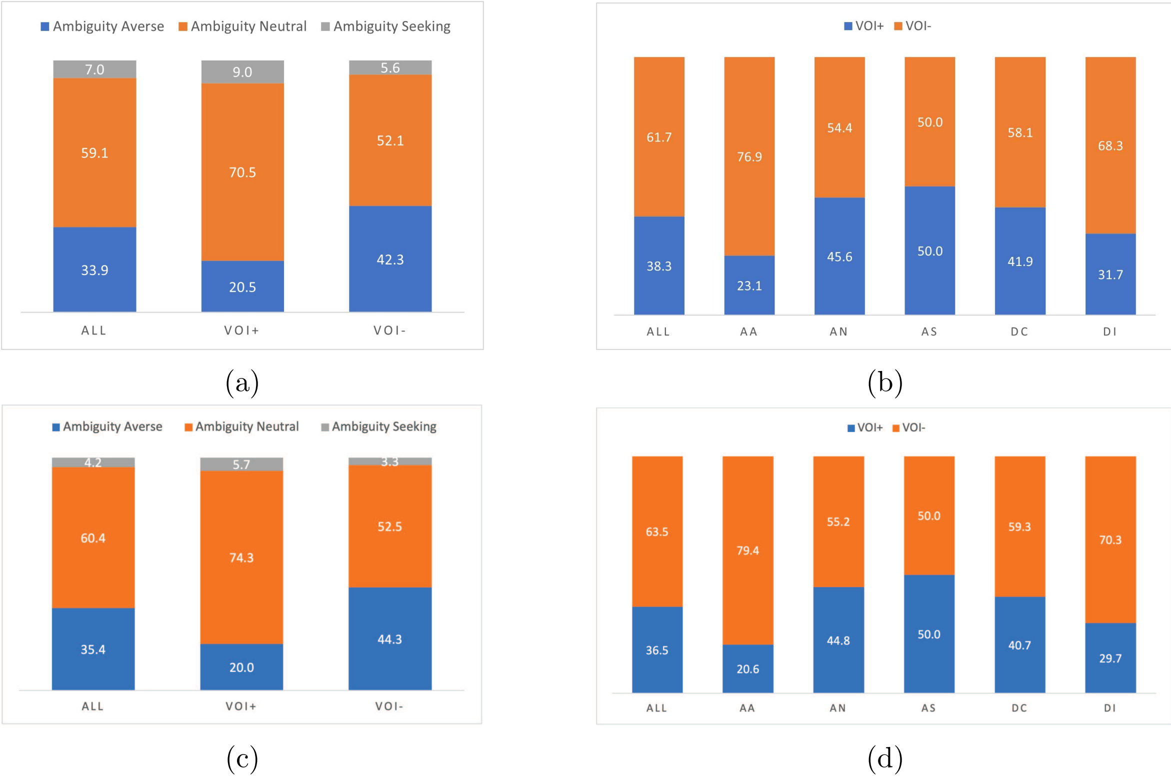

(a) (b) (c) (d)

Figure 5: (a) Distribution of ambiguity attitude for each type of information behavior, (b) distribution of information behavior for each type of ambiguity attitude and dynamic consistency, (c) and (d) show corresponding shares if subjects violating first-order stochastic dominance are excluded (AA ambiguity averse, AN ambiguity neutral, AS ambiguity seeking, DC dynamic consistency, DI dynamic inconsistency)

Thus, our findings are in line with the predictions of recent theoretical contributions. The fact that we were able to replicate the relationship between the VOI and ambiguity attitude put forward by the theory of Li (2020), counteracts worries that our results stem from subjects' choices being random or driven by confusion.

On the other hand, the correlation between VOI− and ambiguity aversion is not perfect. Specifically, 58% of the VOI− group are not ambiguity averse. This suggests that the negative value of information may be more than a mere by-product of ambiguity aversion. Indeed, the drivers for consistent behavior in these two tests may well be of a different nature. To behave in line with the SEU model combined with the standard updating rules in our test of DMs' ambiguity attitudes, that is a variation of the Ellsberg 3-color problem, participants need to be probabilistically sophisticated. On the other hand, to behave like this in our test of the negative value of information, it suffices for participants to ignore the color of the drawn ball.

We asked participants about their confidence in their choices for each choice problem on a 1 to 5 scale (Figure 6a). Participants in the VOI− group were significantly less confident in their choices for Problem 2 than in their choices for Problem 1 compared to participants in the VOI+ group (difference in confidence between decision tasks: VOI− 1.14 vs VOI+ 0.41, $p = 0.002$, two-sided t-test). The larger drop in confidence from Problem 1 to 2 for the VOI− group is exactly what we would expect from participants who prefer to avoid partial information and

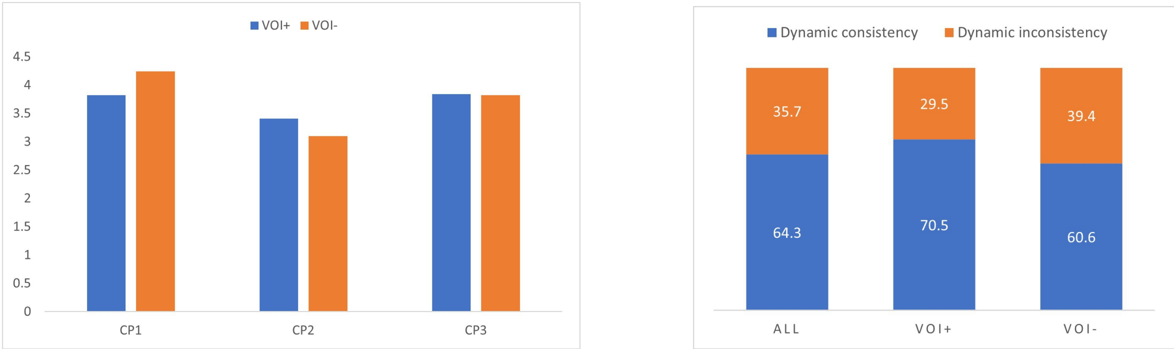

(a) (b)

Figure 6: (a) Average confidence on 1-5 scale for each type of information behavior, Problems 1-3 (b) distribution of dynamic consistency for each type of information behavior

settle for $c_1(\overline{\mathbf{12}}, \overline{\mathbf{34}})$ in Problem 3.

We also asked participants for their estimates of the number of green balls in the Color Info and No Color Info conditions in a questionnaire at the end of the experiment. The average distance between the two estimates was not significantly different for VOI+ and VOI− subjects ($p = 0.208$, two-sided t-test). Therefore, it seems unlikely that the negative value of information was driven by a suspected difference in urn composition.

Finally, in our data, there is no significant correlation between the negative value of information and dynamic inconsistency. Figure 6b shows that the percentage of dynamically inconsistent choices is larger in the VOI− group. But this difference is not statistically significant.

4.1.3 Ambiguity Attitudes

Ambiguity averse subjects constitute 34% of our sample, ambiguity neutral 59%, and ambiguity seeking 7% (Figure 5a). This distribution of ambiguity attitudes is similar to the one reported in Bleichrodt et al. (2021).

We can largely confirm the common finding of a fourfold pattern of ambiguity attitudes (Kocher, Lahno, and Trautmann, 2018) restricted to the gain domain. That is, we observe a moderate fraction of participants, 44%, who prefer the risky option $\overline{\mathbf{1}}$ to the ambiguous $\overline{\mathbf{3}}$ when the likelihood of winning is low, as in Problem 4. Furthermore, when the chances of winning are moderate to high, as in Problem 5, the fraction of participants who prefer the risky option $\overline{\mathbf{34}}$ to the ambiguous option $\overline{\mathbf{14}}$ increases sharply to 83% ($p < 0.001$, two-sided proportion test). As the fourfold pattern suggests, in our experiment, the share of risky choices is significantly higher for high likelihood gains than it is for low likelihood gains.

The proportion of subjects who chose the risky option in Problem 4 and the proportion of subjects who chose the risky option in Problem 5 are both higher in the VOI− group than in the VOI+ group, although not statistically significant (Problem 4: 51% vs 34%, $p = 0.121$; Problem 5: 86% vs 77%, $p = 0.350$; two-sided proportion test). Even this separate analysis of the relationship between ambiguity attitude and the value of information provides further support for the connection between ambiguity aversion and the negative value of information.

4.1.4 Dynamic Consistency

Based on participants' responses to Problems 1 and 2, we can classify them as dynamically consistent or not. Table 10 shows that 64% of all participants in our experiment are dynamically consistent. This number is higher than the 32% reported in Dominiak, Duersch, and Lefort (2012), but very close to the 66% in Bleichrodt et al. (2021). Most of the violations of dynamic consistency, 90%, take the pattern of $\overline{\mathbf{34}}$ in Problem 1 and $\overline{\mathbf{12}}^*$ in Problem 2. The opposite pattern of choosing $\overline{\mathbf{12}}$ and $\overline{\mathbf{34}}^*$ is responsible for only 10% of the overall violations, which is not surprising, given that $\overline{\mathbf{12}}$ is the stochastically dominated option in Problem 1.

Table 10: Share of dynamically consistent choices

Problem 1 & 2	Share	N
$\overline{\mathbf{12}} \succ \overline{\mathbf{34}}$ iff $\overline{\mathbf{12}}^* \succ \overline{\mathbf{34}}^*$	64% DC ***	115

Notes: DC = dynamic consistency. Column "Share" shows percentage of choices $\overline{\mathbf{12}}$ and $\overline{\mathbf{12}}^*$ (resp., $\overline{\mathbf{34}}$ and $\overline{\mathbf{34}}^*$) in Problems 1 and 2. Binary choice: Two-sided binomial test against $p = 0.5$. *** denotes significance at 1%.

The result changes when we look at the subgroup of participants who are ambiguity averse according to the last two choice problems. In line with the results by Dominiak, Duersch, and Lefort (2012), ambiguity averse participants are significantly more often dynamically inconsistent (51% vs. 28%, $p = 0.021$, two-sided proportion test). On the other hand, ambiguity neutral participants are more often dynamically consistent (75% vs. 49%, $p = 0.008$, two-sided proportion test).

4.1.5 Gender and Age

Finally, we find a weak gender effect insofar as the share of females is larger in the VOI− group than in the VOI+ group (63% vs 46%, $p = 0.091$, two-sided proportion test). Furthermore, the percentage of VOI− among females is 18% larger than this percentage is among males (69% vs 51%, $p = 0.075$, two-sided proportion test).

Finally, we also find that the average age in the VOI− group is lower than in the VOI+ group (22.5 vs 24.4, $p = 0.125$, two-sided t-test).

See Appendix B for a table listing all correlations.

4.2 Results of Second Experiment

4.2.1 Construction of Variables

Each of the ten choice problems in the second experiment consists of 15 lines, where the top 7 lines offer extra payments for the choice of the first option, while the bottom 7 lines pay extra for the choice of the second option. The extra payments are €5.00, €2.00, €1.50, €1.00, €0.75, €0.50, and €0.25. The middle (8th) line offers a choice between the two options with no extra payments, i.e., the plain urn-bets of the corresponding choice problem.

For each choice problem, define the value as the extra payment added to the first option that makes the DM indifferent between the two. Extra payments added to the second option are considered as negative numbers. Assuming no multiple switching, the elicited value is the mid-point between the two extra payments where the switch occurred.[6]

VOI is defined as the value in Problem 5, which is an extra payment added to the no-information bet that makes the DM indifferent between this bet and the information bet. For the complexity and ambiguous reliability tests, VOI is defined similarly as the value in Problem 8 and Problem 10 respectively. Note that the elicited value of ± 0.125 does not contradict the hypothesis that a participant was indifferent between the two options and was forced to break the indifference when answering the middle line of the list. To deal with this problem, we use a conservative approach and classify values ± 0.125 as VOI zero rather than positive or negative, which may overstate this category.

Let x_i be the value in Problem i as defined above. An ambiguity averse subject prefers the first option in the middle line of Problem 1 (so that $x_1 < 0$) and the second option in the middle line of Problem 2 ($x_2 > 0$). The ambiguity premium is defined as $x_2 - x_1$ if x_1 and x_2 have different signs and zero otherwise. Thus, the ambiguity premium is positive for ambiguity averse subjects, zero for neutral, and negative for seeking. Analogously, the dynamic consistency premium is defined as $x_3 - x_4$ if x_3 and x_4 have different signs and zero otherwise. It is (non-)zero for dynamically (in-)consistent subjects.[7]

4.2.2 Summary Statistics

107 subjects participated in the second experiment. The average age was 21.2 years, the share of economics students was 23%, and the share of female participants was 68%. 7 subjects

[6] For example, if a participant chose A plus €0.75 against B, but B against A plus €0.5, the elicited value is 0.625.

[7] The ambiguity premium of ± 0.25 is compatible with ambiguity neutrality because of the possible indifference. We classified such subjects as ambiguity neutral. Similarly, the dynamic consistency premium ± 0.25 is compatible with the dynamically consistent behavior and is classified as such.

did multiple switching in at least one choice problem. Since it is unclear how to define the corresponding values in these cases, we excluded these subjects and analyzed the remaining 100 observations.

4.2.3 Negative Value of Information

Our conservative approach to the data of the second experiment **identifies 35% of the participants as showing a negative value of information.**[8] This is lower than in the first experiment (62%). Other than by the conservative approach, this difference may also be explained by a psychological bias toward the middle of a symmetric multiple price list (Andersen et al., 2006; List, 2007), resulting in our experiment in $VOI = 0$. Zooming in on the negative VOI, it lies between 0 and -0.5 for 16%, between -0.5 and -1 for 17%, and below -1 for 2% of participants.

As for the link between the negative VOI and ambiguity aversion, **the second experiment showed a strong connection between the two.** Table 11 shows that the percentage of ambiguity aversion in the negative VOI group (80%) is much higher than in the zero VOI and positive VOI groups (11% and 14% resp.). The difference in the proportion of ambiguity aversion between negative VOI and non-negative VOI subjects is significant (80% vs 12%, $p < 0.001$, two-sided proportion test), as is the difference in the mean ambiguity premium (1.21 vs -0.11, $p < 0.001$, two-sided t-test). An OLS regression gives a negative relation between VOI and the ambiguity premium ($t = -5.29$, $p < 0.001$).[9] The likelihood of having a negative VOI is positively related to the ambiguity premium ($z = 3.08$, $p = 0.002$, probit).

Table 11: Distribution of ambiguity attitude and dynamic consistency for each type of information behavior

	all	negative	zero	positive
Ambiguity Attitude				
mean ambiguity premium	0.36	1.21	-0.08	-0.14
median ambiguity premium	0	1	0	0
averse	36%	80%	11%	14%
neutral	56%	20%	81%	69%
seeking	8%	0%	8%	17%
Dynamic Consistency				
mean DC premium	0.19	0.25	0.15	0.17
median DC premium	0	0	0	0
consistent	82%	77%	92%	76%
inconsistent	18%	23%	8%	24%

Notes: Columns “negative”, “zero”, and “positive” refer to VOI categories.

[8] This approach identifies as $VOI = 0$ all cases where indifference cannot be ruled out. Assuming no indifference, the share of the negative value of information is 59%.

[9] We eliminated 5 outliers ($N = 95$).

As Table 12 shows, 78% of ambiguity averse subjects showed a negative VOI. For 52% of ambiguity neutral subjects VOI was zero, while the majority of ambiguity seeking subjects (62%) had a positive VOI. Clearly, there is a difference in the mean VOI for the ambiguity averse and non-averse subjects (-0.37 vs. 0.32, $p < 0.001$, two-sided t-test).

There is no relation between VOI and DC premium ($t = -0.32$, $p = 0.75$, OLS regression after elimination of outliers). Although Table 11 shows that the proportion of dynamic consistency among the VOI zero subjects is higher than among the rest (92% vs. 77%, $p = 0.1$, two-sided proportion test), no strong conclusion in this respect can be drawn from the data.

While there were weak gender effects in the first experiment, we did not find any significant gender or age effects in the second experiment.

Table 12: Distribution of VOI for each type of ambiguity attitude and dynamic consistency

	all	averse	neutral	seeking	consistent	inconsistent
			VOI - ambiguity			
mean	0.07	-0.37	0.21	1.08	0	0.42
median	-0.12	-0.38	0.12	0.88	-0.12	-0.12
negative	35%	78%	12%	0%	33%	44%
zero	36%	11%	52%	38%	44%	17%
positive	29%	11%	36%	62%	27%	39%
			VOI - complexity			
mean	0.39	0.35	0.3	1.2	0.32	0.69
median	0.12	0.38	0.12	0.75	0.12	0.25
negative	8%	8%	9%	0%	5%	22%
zero	46%	36%	55%	25%	50%	28%
positive	46%	56%	36%	75%	45%	50%
			VOI - ambiguous reliability			
mean	-0.01	-0.12	0.03	0.23	-0.04	0.13
median	0.12	-0.12	0.12	0.12	0.12	-0.12
negative	13%	11%	12%	25%	12%	17%
zero	75%	81%	79%	25%	77%	67%
positive	12%	8%	9%	50%	11%	17%

Notes: Columns “averse”, “neutral”, and “seeking” refer to ambiguity attitudes. The last two columns are dynamic consistency classes. Percentages need not always add up to 100% due to rounding.

4.2.4 Ambiguity Attitudes, Dynamic Consistency, and Beliefs

In the second experiment, the distribution of ambiguity attitudes was 36% averse, 56% neutral, and 8% seeking, which is very close to the results of the first experiment. At 82%, the share of dynamic consistency was higher than in the first experiment.

In half of the sessions, the participants were asked for their estimates of the number of green balls via a questionnaire at the beginning of the experiment. This could induce them to form beliefs about the composition of the urns. Their answers to the subsequent choice problems were not significantly different from those who received this questionnaire at the end of the experiment.[10] This suggests that subjects form beliefs about the urn composition independent of the questionnaire.[11]

4.2.5 Complexity

In both experiments, the problems where the color of the drawn ball is known up front are arguably more complex than those where the color is not known up front. This suggests that the negative VOI may also be driven by complexity avoidance (Bruce and Johnson, 1996; Huck and Weizsäcker, 1999; Bayer and Ke, 2013; Duttle and Inukai, 2015; Zilker, Hertwig, and Pachur, 2020) rather than by ambiguity aversion. To test this explanation, we measured VOI in a situation which is almost identical to our standard test of the negative value of information. The only difference is that the participants are now informed about how the exact number n_3 of balls marked with a 3 in the urns is determined. The number n_3 is determined randomly such that every number 0, 1, 2,..., 11 is equally likely for n_3 (and n_4 is defined as $11 - n_3$). As such, the situation here does not involve ambiguity but retains the complexity associated with the information about the color of the drawn ball.

As the second section of Table 12 shows, **only 8% of subjects showed a negative VOI in this situation** (vs. 35% in the ambiguous situation, $p < 0.001$, two-sided proportion test). As expected, this number is practically the same for ambiguity averse (8%) and neutral (9%) subjects, so is the mean VOI (0.35 vs. 0.3, $p = 0.7$, two-sided t-test). There is also no connection between VOI and the ambiguity premium in this case ($t = -0.91$, $p = 0.36$, OLS regression).

The conclusion we can draw from our data is that complexity may influence the results, but its effect is not sufficient to explain the findings.

4.2.6 Information of Ambiguous Reliability

Shishkin and Ortoleva (2023) found no correlation between the negative VOI and ambiguity aversion. One notable difference to our experiment here lies in the mechanism which determines how information generates ambiguity. In Shishkin and Ortoleva (2023), the set of beliefs dilates because the information has ambiguous reliability. When we make the information about the ball color less reliable in our experiment, the correlation between the negative VOI and ambiguity aversion vanishes, as well.

[10] There is no significant difference in the mean VOI, the proportion of the negative VOI, and the mean ambiguity premium.

[11] We would like to thank an anonymous referee for suggesting this test to us.

The last section of Table 12 shows that **only 13% of subjects demonstrated a negative VOI in the situation when information is of ambiguous reliability.** VOI was zero for the majority of subjects (75%), similarly to Shishkin and Ortoleva (2023). Moreover, the percentage of negative VOI is similar for the ambiguity averse and neutral subjects, and there is no significant difference in mean VOI. There is no strong connection between VOI and ambiguity premium ($t = -1.94$, $p = 0.056$, OLS regression) either, also when restricted to ambiguity averse subjects ($t = -0.68$, $p = 0.5$, OLS regression). There is also no connection between the likelihood of having a non-zero VOI and ambiguity premium ($z = -0.62$, $p = 0.54$, probit).

Therefore, the reliability of the information plays a key role. When ambiguity is generated by information of ambiguous reliability, the negative VOI and its connection to ambiguity aversion disappear.

4.2.7 Pooled Data and Robustness Check

Table 13 shows the results from the combined data of our two experiments. As a robustness check, we excluded from the analysis subjects who chose a stochastically dominated option in any of the questions, gave inadequate estimates when asked about the composition of the urns in the questionnaire, or reported minimal strength of preference estimates in the first experiment that could be interpreted as indifference. Table 13 shows that the main findings from our experiment hold for both the main and the reduced sample.

Table 13: Pooled data

Only	$VOI < 0$	AA in $VOI \geq 0$	AA in $VOI < 0$	p-value	N
all	49%	16%	63%	< 0.001	215
robust	49%	15%	63%	< 0.001	169

Notes: Column "$VOI < 0$" shows the proportion of the negative VOI. Next two columns compare percentages of ambiguity averse choices among those with the negative VOI and the rest (two-sided proportion test).

5 Discussion

This section provides a detailed discussion of the assumptions underlying our experiment and the specific design choices we made.

5.1 Underlying Assumptions

Assumptions 1-3 were necessary to keep the experiment brief and reduce the number of choice problems that participants have to answer, and to simplify the subsequent analysis.

First of all, Assumption 1 implies that the DM holds similar beliefs about the compositions of the urns in the Color Info and No Color Info conditions. This is necessary for claiming that $\overline{\mathbf{12}}$ and $\overline{\mathbf{12}}^*$ (correspondingly, $\overline{\mathbf{34}}$ and $\overline{\mathbf{34}}^*$) are essentially the same act shifted by €0.5. We control for its validity in a questionnaire at the end of the experiment. While the assumption may not hold in general, we found no evidence that the difference in beliefs about urn compositions drives our result as Section 4.1.2 establishes.

Translation invariance follows from Assumption 1 and implies that if the payoffs of two acts are increased or decreased by the same fixed amount in each state, the preference between these two acts does not change. This allows us to derive the preference between $\overline{\mathbf{12}}^*$ and $\overline{\mathbf{34}}^*$ when the color of the drawn ball is unknown from the preference between $\overline{\mathbf{12}}$ and $\overline{\mathbf{34}}$. Recall that we only asked about the preference between $\overline{\mathbf{12}}^*$ and $\overline{\mathbf{34}}^*$ when the color of the drawn ball was known. Our test of dynamic consistency relies on this property to hold. For a relatively low increase in payoffs as is the case in our experiment (€0.5), preferences presumably satisfy translation invariance. For large payoff increases this assumption may become problematic, as was recently shown by Baillon and Placido (2019) and König-Kersting, Kops, and Trautmann (2023). Translation invariance is satisfied in SEU, MEU, and α-MEU.

Symmetry arguments play an important role throughout the paper. Assumptions 2 and 3 allow us not to care about which of the two colors, green or blue, realizes in an experimental session. For the MEU model, these assumptions would be satisfied by a set of beliefs which is symmetric with respect to 1 and 2, and, respectively, 3 and 4. Given the information about the urn composition available to the participants, this seems plausible if color preferences are neglected.

Finally, it might not be obvious how the DM aggregates the conditional utilities V_A and $V_{\bar{A}}$ when she estimates the ex-ante value of information $\{A, \bar{A}\}$ under ambiguity. For Definition 2.2, we used the rule axiomatized in Li (2020). However, our results regarding the sign of the value of information would remain valid under milder assumptions, as well. For our design, it is sufficient to assume that if both $V_A < V_0$ and $V_{\bar{A}} < V_0$, then the value of information is negative. In other words, if the DM is worse off in any of the possible scenarios that information offers, then the value of such information is negative. Hence, our design works for any aggregation rule satisfying this property.

5.2 Critical Discussion of Design Choices

This section aims to provide a critical discussion of some of our design choices.

First, the directness of our test. The primitive concept in economics for modeling individual behavior, beliefs, and attitudes is choice. Rather than information structures themselves, the traditional objects of choice in decision theory are sure alternatives, lotteries, or, acts. Respecting this tradition, the formal definition of the value of information defines this value in an indirect way by resting on the values a DM assigns to an act and its conditional counterparts. We follow this tradition with our test of information aversion that is based on choices between an act and its conditional counterparts. Another way to test whether the value of information is negative would be to directly ask participants whether they want to receive additional information about the state of the world or not.

None of these approaches is without its drawbacks. While our design rests on the assumptions in Section 5.1 to hold, under the alternative design participants would have to imagine the more informed situation, they would not be able to experience both situations, but would still be required to make a decision between the two. Furthermore, this alternative design would be susceptible to default effects that, for tests of strategic ignorance, Grossman (2014) has shown to play an important role when participants are asked to choose between different information structures. Specifically, their analysis shows that the frequency of information-averse choices ranges between 3% and 45% depending on whether ignorance is the default option or not.

Second, the random incentive mechanism. Participants in our experiment make decisions in several choice problems and only one of these decisions is randomly selected to be payoff-relevant at the end of the experiment. This random incentive mechanism can become an issue if participants do not consider the decision situations in isolation (Holt, 1986; Karni and Safra, 1987; Freeman, Halevy, and Kneeland, 2019; Baillon, Halevy, and Chen, 2022b). The issue then lies in the measurement of ambiguity attitudes being confounded (Baillon, Halevy, and Chen, 2022a). Oechssler and Roomets (2014) and Bade (2015) were among the first to point out that hedging motives may lead to a misclassification of ambiguity averse participants as subjective expected utility maximizers. Translating their argument into our setup, would mean that participants use the choice problems to bet on and against the same ball-numbers to arrive at an unambiguous expected value. The empirically observed patterns speak against this interpretation. The hedging explanation does not predict the strong connection between negative VOI and ambiguity aversion ("AA in $VOI \geq 0$" vs. "AA in $VOI < 0$": 15% vs. 63%, $p < 0.001$, $n = 169$, two-sided test of proportions, reduced pooled data). In the setting of our experiment, we would expect hedging effects —if present and sizable— to result in a weaker such connection.

Third, information aversion may be driven by quantitative sophistication or a lack thereof

(Abdellaoui, Klibanoff, and Placido, 2015). To tackle this potential confound, we may use whether subjects stated in the questionnaire at the end of our experiment that they took a statistics class or that they are economics students. As for the statistics class, there is no significant difference between subjects who took a statistics class vs. those who did not, when it comes to negative value of information (60% vs. 63%, $p = 0.861$, two-sided proportion test, first experiment). Likewise, there is no significant difference between subjects who are economics students vs. those who are not (55% vs. 64%, $p = 0.535$, two-sided proportion test, first experiment).

Fourth, multiple switching. This is not an issue for our first experiment on the negative value of information, since it does not rely on MPLs. Our second experiment, however, does rely on them. The MPLs used in our second experiment may induce participants to switch back and fourth between the two bets in a choice problem when the guaranteed extra payments vary from high to low. Typical studies in developed countries find multiple switching to affect approximately 10% of subjects (Holt and Laury, 2002; Dave et al., 2010). To reduce the occurrence of multiple switching, we followed Bruner (2011) and emphasized that only one question will determine earnings, before decision sheets were distributed. In line with the multiple-switching frequencies of $2.3\% - 6.7\%$ reported by Bruner (2011), for our second experiment this resulted in a total of 6.5% (7 out of 107) of participants to switch more than once between the two columns of the MPL of at least one choice problem.

Fifth, complexity avoidance. Halevy (2007) and numerous subsequent contributions found that individuals perceive compound risk similarly to ambiguity. One potential explanation for this may lie in decision situations with ambiguity and compound risk sharing comparable complexity. This is the idea behind the part of the second experiment that aims at disentangling the effects of complexity avoidance and ambiguity aversion. The results suggest that when it comes to information, the difference in the perception of compound risk and ambiguity is significant and warrants further investigation.

6 Conclusion

The findings from our experiments show that, at 49%, the percentage of individuals showing a negative value of information is substantial and robust. As such, the value of information can indeed be negative when new information renders hedging against ambiguity impossible. Note that under SEU, individuals are "free to ignore the observation. That obvious fact is the theory's expression of the commonplace that knowledge is not disadvantageous" (Savage, 1954, p. 107). Our results can, thus, be interpreted as evidence that many individuals are unable to do so. Moreover, the value of information is correlated with ambiguity aversion. Neither complexity avoidance nor information with ambiguous reliability can reproduce the results.

While the economics and psychology literature puts forward reasons for information aversion such as disappointment aversion, regret aversion, optimism maintenance and belief investments (Golman, Hagmann, and Loewenstein, 2017), ambiguity theories predict that new information can have a negative value when it breaks an existing hedge against ambiguity. Our experiment provides support for this connection and the ambiguity theories predicting it (Li, 2020).

Our results suggest that sometimes less information may be perceived as a better option. Note in this direction, that women at risk of developing breast cancer delay visiting the doctor (Caplan, 1995; Meechan, Collins, and Petrie, 2002; Kőszegi, 2003), investors reject free and independent investment advice (Lacko and Pappalardo, 2010; Bhattacharya et al., 2012; Kosfeld and Schüwer, 2017), and financial education has a small impact on decision-making (Choi, Laibson, and Madrian, 2010). Our results suggest that counseling on topics like health, financial literacy and others may have to be delivered in a way that takes into account what effect the information has on the subjectively perceived ambiguity in the situation.

Appendix A

In this appendix, we analyze Problems 1-5 of the first experiment using prominent decision-theoretic models.

Let the choice set be $F = \{\overline{\mathbf{12}}, \overline{\mathbf{34}}\}$ as in Problem 1. In the SEU model, consider the value of information $\{G, B\}$ defined by Equation (3). Assume $p(1) = \frac{5}{21}$, $p(2) = \frac{5}{21}$, and the uniform belief $p(3) = p(4) = \frac{5.5}{21}$ on $\{3, 4\}$ following the principle of insufficient reason (Eichberger and Pasichnichenko, 2021). Then, all three maxima V_G, V_B, and V_0 are attained at $\overline{\mathbf{34}}$. Clearly, the information $\{G, B\}$ has zero value in the SEU model. Hence, such a DM would not forego the increase in payments offered by Problem 2.

To apply MEU to our experiment, assume that the DM's set of beliefs over $S = \{1, 2, 3, 4\}$ is given by

$$p(1) = \frac{5}{21}, \quad p(2) = \frac{5}{21}, \quad p(3) = \frac{5.5 + x}{21}, \quad p(4) = \frac{5.5 - x}{21}, \quad -\varepsilon \leq x \leq \varepsilon \tag{8}$$

where $0 \leq \varepsilon \leq 5.5$ measures the DM's subjective ambiguity. In other words, the DM thinks that the number n_3 of 3-balls is in the range between $5.5 - \varepsilon$ and $5.5 + \varepsilon$. For instance, if $\varepsilon = 2.5$, then the DM thinks that the urn contains between three and eight balls marked with "3". Note that this set of beliefs is not rectangular (Epstein and Schneider, 2003). In principle, subjects may form a set of subjective beliefs that includes $\mathcal{C}$ and is rectangular (Riedel, Tallon, and Vergopoulos, 2018). In this case, however, the value of information is zero (Li, 2020). Since in our experiment most subjects show a negative VOI, we abstain from assuming rectangularity here.

Without loss of generality, we can normalize the utility of money u such that $u(4) = 0$ and $u(10) = 1$. Then, for Problem 1, we have the minimal expected utility $\frac{10}{21}$ for $\overline{\mathbf{12}}$ and $\frac{11}{21}$ for $\overline{\mathbf{34}}$, so that $\overline{\mathbf{34}} \succ \overline{\mathbf{12}}$ and $V_0 = \frac{11}{21}$. Suppose now that while choosing between $\overline{\mathbf{12}}$ and $\overline{\mathbf{34}}$, the DM is offered the information about the color of the drawn ball. By symmetry, $V_B = V_G$, so that Equation (6) reduces to

$$VOI = V_G - V_0$$

and we only have to find V_G. The Full Bayesian update $\mathcal{C}_G$ is given by

$$p(1) = \frac{5}{10.5 + x}, \quad p(2) = 0, \quad p(3) = \frac{5.5 + x}{10.5 + x}, \quad p(4) = 0, \quad -\varepsilon \leq x \leq \varepsilon.$$

Then, for the conditional maxmin expected utility M_G of $\overline{\mathbf{12}}$ and $\overline{\mathbf{34}}$, we have

$$M_G(\overline{\mathbf{12}}) = \frac{5}{10.5 + \varepsilon}, \quad M_G(\overline{\mathbf{34}}) = \frac{5.5 - \varepsilon}{10.5 - \varepsilon}.$$

Thus, $\overline{\mathbf{12}} \succcurlyeq_G \overline{\mathbf{34}}$ for $\varepsilon \geq \varepsilon^* = \frac{\sqrt{21}}{2} \approx 2.3$ and $\overline{\mathbf{34}} \succ_G \overline{\mathbf{12}}$ otherwise, which implies $V_G = M_G(\overline{\mathbf{12}})$

for $\varepsilon \geq \varepsilon^*$ and $V_G = M_G(\mathbf{34})$ otherwise. Simple arithmetic computations show that $V_G < \frac{11}{21} = V_0$ for any $0 < \varepsilon \leq 5.5$, so that $VOI < 0$. Only in the extreme case where $\mathcal{C}$ is a singleton set, we have $\varepsilon = 0$ and $VOI = 0$. Therefore, for an ambiguity averse DM with MEU-preferences, the information about the color of the drawn ball has a negative value.

Figure 7: The value of the information $\{G, B\}$ under MEU-preferences and in relation to ambiguity (epsilon)

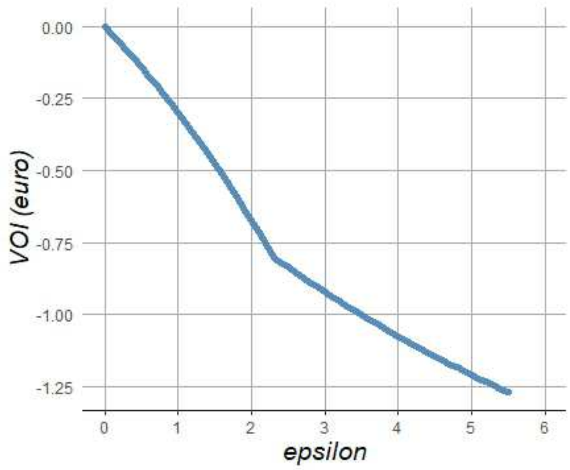

Greater ambiguity corresponds to more money that the DM will pay/forego to avoid the information $\{G, B\}$. Figure 7 shows this relationship assuming a linear utility function of money. Put differently, it illustrates the monetary amount required to compensate for the negative value of the information $\{G, B\}$. For example, if the DM thinks that the urn contains between three and eight 3-balls, which corresponds to $\varepsilon = 2.5$, then she will prefer the Color Info condition to the No Color Info condition only if the bonus payment is at least €0.84.

Note that the MEU model can accommodate only ambiguity aversion (and neutrality when the set of beliefs is a singleton). Since 66% of participants in the first experiment were ambiguity neutral or seeking, we also consider the extension of α-MEU (Ghirardato, Maccheroni, and Marinacci, 2004). In this model, the DM ranks acts by maximizing a weighted average of the minimal and maximal expected utility of act f with respect to the set of beliefs. Here, the value of information $\{A, \bar{A}\}$ is defined as

$$VOI = \alpha \min_{p \in \mathcal{C}} \left[p(A)V_A + p(\bar{A})V_{\bar{A}}\right] + (1-\alpha) \max_{p \in \mathcal{C}} \left[p(A)V_A + p(\bar{A})V_{\bar{A}}\right] - V_0,$$

where

$$V_E = \max_{f \in F} \left[\alpha \min_{p_E \in \mathcal{C}_E} \sum_{s \in E} p_E(s) u\left(f(s)\right) + (1-\alpha) \max_{p_E \in \mathcal{C}_E} \sum_{s \in E} p_E(s) u\left(f(s)\right)\right]$$

and $\mathcal{C}_E$ is the Full Bayesian update of $\mathcal{C}$. Since, in our experiment, $V_G = V_B$ by symmetry, the expression for VOI collapses to $VOI = V_G - V_0$.

Figure 8a summarizes the results of numerically modeling this VOI, for all possible values of ε and α, under the assumption that the set of beliefs is defined by Equation (8). Figure 8b

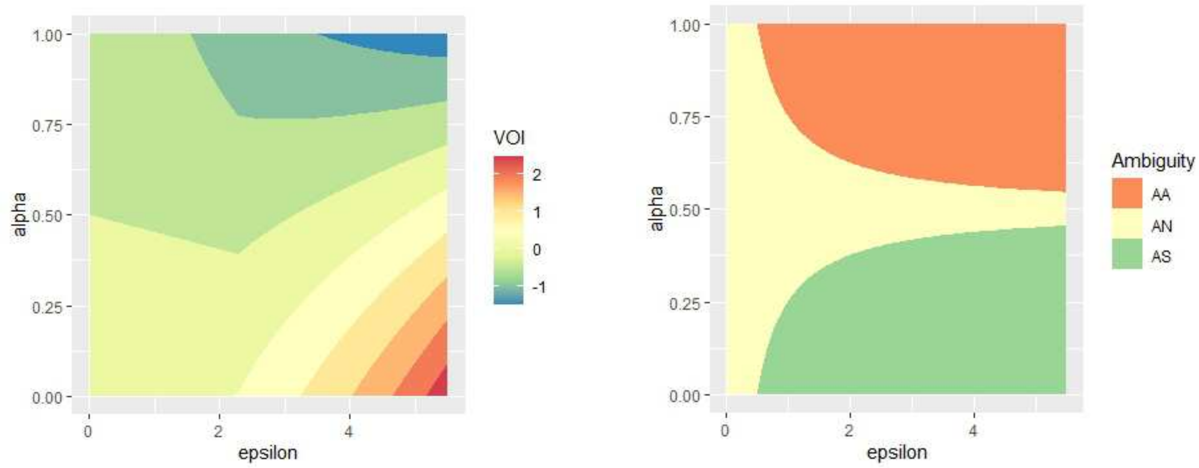

(a) (b)

Figure 8: (a) The predicted value of information $\{G, B\}$ for different pairs of (ε, α) in α-MEU. One grade of the color scale corresponds to a €0.5 change. $VOI < 0$ for the three upper right segments. $VOI < 0.5$ for the two upper right segments. (b) The predicted ambiguity attitudes in Problems 4-5 for different pairs of (ε, α) in α-MEU (AA averse, AN neutral, AS seeking).

shows how different DMs with α-MEU preferences would perform in our test of their ambiguity attitude. This shows that, for α-MEU, the value of information can only go below €0.5 for ambiguity averse subjects. In our experiment, however, the share of ambiguity aversion with a negative value of information is a mere 42%, which suggests that the α-MEU model points in the right direction, but may be incomplete.

Segal (1990) views ambiguity as a two-stage lottery where the first (imaginary) stage is over all possible values of the probability of the events in the second stage. Ambiguity attitudes are represented by a rank-dependent probability weighting function in the first stage. In this model, DMs can switch from ambiguity aversion to ambiguity seeking when the likelihood gets smaller in the way the fourfold pattern of ambiguity attitudes in Kocher, Lahno, and Trautmann (2018) suggests it. In the interest of brevity and since we only found a mere 7% of subjects being ambiguity seeking, we leave the calculations for the interested reader.

While SEU-preferences combined with the standard updating rules are always dynamically consistent, note that in the MEU model, we have $\overline{\mathbf{34}} \succ \overline{\mathbf{12}}$, but $\overline{\mathbf{12}} \succ_G \overline{\mathbf{34}}$ and, by symmetry, $\overline{\mathbf{12}} \succ_B \overline{\mathbf{34}}$ for $\varepsilon \geq \varepsilon^*$, which is dynamically inconsistent.

Observe that SEU-preferences are always ambiguity neutral, while MEU-preferences can be ambiguity averse. To see this, consider Problem 4. After normalizing utility, the minimal expected utility M is

$$M(\overline{\mathbf{1}}) = \frac{5}{21}, \quad M(\overline{\mathbf{3}}) = \frac{5.5 - \varepsilon}{21},$$

so that $\overline{\mathbf{1}} \succ \overline{\mathbf{3}}$ if $\varepsilon > 0.5$. On the other hand, in Problem 5, we have

$$M(\overline{\mathbf{14}}) = \frac{10.5 - \varepsilon}{21}, \quad M(\overline{\mathbf{34}}) = \frac{11}{21}$$

which implies $\overline{\mathbf{34}} \succ \overline{\mathbf{14}}$ for any such ε. Therefore, any DM with MEU-preferences and $\varepsilon > 0.5$ is classified as ambiguity averse in our experiment.

Appendix B

Table B1: Correlations in the first experiment

	VOINEG	AGE	RELIG	ECON	STAT	RIGHT	MALE	DYNINC	AMBA	AMBS
AGE	-0.17*									
RELIG	-0.04	0.06								
ECON	-0.08	-0.20**	0.01							
STAT	-0.03	-0.04	0.05	0.25***						
RIGHT	-0.11	-0.07	0.09	0.00	0.01					
MALE	-0.19**	0.13	-0.05	0.15	0.27***	0.17*				
DYNINC	0.10	-0.08	0.03	-0.10	-0.09	-0.27***	-0.09			
AMBA	0.22**	-0.10	0.02	-0.12	0.09	-0.12	-0.10	0.23**		
AMBS	-0.07	-0.02	-0.01	0.00	-0.26***	0.06	0.04	0.08	-0.20**	
AMBN	-0.18*	0.10	-0.01	0.12	0.05	0.09	0.07	-0.27***	-0.86****	-0.33****

Notes: $N = 115$. VOINEG = negative value of information, AGE = participant's age, RELIG = religious, ECON = economics student, STAT = took a statistics course, RIGHT = right political views, MALE = male gender, DYNINC = dynamically inconsistent, AMBA = ambiguity averse, AMBS = ambiguity seeking, AMBN = ambiguity neutral. Two-sided t-test. * denotes significance at 10%, ** at 5%, *** at 1%, and **** at 0.1%.

References

Abdellaoui, Mohammed, Peter Klibanoff, and Lætitia Placido. 2015. "Experiments on Compound Risk in Relation to Simple Risk and to Ambiguity." *Management Science* 61 (6):1306–1322.

Andersen, Steffen, Glenn W Harrison, Morten Igel Lau, and E Elisabet Rutström. 2006. "Elicitation Using Multiple Price List Formats." *Experimental Economics* 9:383–405.

Andries, Marianne and Valentin Haddad. 2020. "Information Aversion." *Journal of Political Economy* 128 (5):1901–1939.

Aumann, Robert J. 1974. "Subjectivity and Correlation in Randomized Strategies." *Journal of Mathematical Economics* 1 (1):67–96.

Bade, Sophie. 2015. "Randomization Devices and the Elicitation of Ambiguity-Averse Preferences." *Journal of Economic Theory* 159:221–235.

Baillon, Aurélien, Yoram Halevy, and Li Chen. 2022a. "Experimental Elicitation of Ambiguity Attitude using the Random Incentive System." *Experimental Economics* 25:1002–1023.

———. 2022b. "Randomize at your own Risk: On the Observability of Ambiguity Aversion." *Econometrica* 90:1085–1107.

Baillon, Aurélien and Lætitia Placido. 2019. "Testing Constant Absolute and Relative Ambiguity Aversion." *Journal of Economic Theory* 181:309–332.

Bayer, Ralph-C. and Changxia Ke. 2013. "Discounts and Consumer Search Behavior: The Role of Framing." *Journal of Economic Psychology* 39:215–224.

Bhattacharya, Utpal, Andreas Hackethal, Simon Kaesler, Benjamin Loos, and Steffen Meyer. 2012. "Is Unbiased Financial Advice to Retail Investors Sufficient? Answers from a Large Field Study." *The Review of Financial Studies* 25 (4):975–1032.

Blackwell, David. 1953. "Equivalent Comparisons of Experiments." *The Annals of Mathematical Statistics* 24 (21):265–272.

Bleichrodt, Han, Jürgen Eichberger, Simon Grant, David Kelsey, and Chen Li. 2021. "Testing Dynamic Consistency and Consequentialism under Ambiguity." *European Economic Review* 134:103687.

Brown, Gordon D. A. and Lukasz Walasek. 2020. "Models of Deliberate Ignorance in Individual Choice." In *Deliberate Ignorance: Choosing not to Know*, edited by R. Hertwig and C. Engel. Cambridge, MA: MIT Press, 135–151.

Bruce, A. C. and J. E. V. Johnson. 1996. "Decision-Making under Risk: Effect of Complexity on Performance." *Psychological Reports* 79 (1):67–76.

Bruner, David M. 2011. "Multiple Switching Behaviour in Multiple Price Lists." *Applied Economics Letters* 18 (5):417–420.

Brunnermeier, Markus K and Jonathan A Parker. 2005. "Optimal Expectations." *American Economic Review* 95 (4):1092–1118.

Caplan, Lee S. 1995. "Patient Delay in Seeking Help for Potential Breast Cancer." *Public Health Reviews* 23 (3):263–274.

Choi, James J, David Laibson, and Brigitte C Madrian. 2010. "Why Does the Law of one Price Fail? An Experiment on Index Mutual Funds." *The Review of Financial Studies* 23 (4):1405–1432.

Cohen, Michele, Itzhak Gilboa, Jean-Yves Jaffray, and David Schmeidler. 2000. "An Experimental Study of Updating Ambiguous Beliefs." *Risk, Decision and Policy* 5 (2):123–133.

Dana, Jason, Roberto A Weber, and Jason Xi Kuang. 2007. "Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness." *Economic Theory* 33 (1):67–80.

Dave, Chetan, Catherine C Eckel, Cathleen A Johnson, and Christian Rojas. 2010. "Eliciting Risk Preferences: When is Simple Better?" *Journal of Risk and Uncertainty* 41:219–243.

Deshpande, Rohit and Ajay K Kohli. 1989. "Knowledge Disavowal: Structural Determinants of Information-Processing Breakdown in Organizations." *Science Communication* 11 (2):155–169.

Dillenberger, David. 2010. "Preferences for one-shot Resolution of Uncertainty and Allais-type Behavior." *Econometrica* 78 (6):1973–2004.

Dominiak, Adam, Peter Duersch, and Jean-Philippe Lefort. 2012. "A Dynamic Ellsberg Urn Experiment." *Games and Economic Behavior* 75 (2):625–638.

Duttle, Kai and Keigo Inukai. 2015. "Complexity Aversion: Influences of Cognitive Abilities, Culture and System of Thought." *Economics Bulletin* 35:846–855.

Eichberger, Jürgen and Illia Pasichnichenko. 2021. "Decision-Making with Partial Information." *Journal of Economic Theory* 198:105369.

Eliaz, Kfir and Ran Spiegler. 2006. "Can Anticipatory Feelings Explain Anomalous Choices of Information Sources?" *Games and Economic Behavior* 56 (1):87–104.

Ellsberg, Daniel. 1961. "Risk, Ambiguity, and the Savage Axioms." *The Quarterly Journal of Economics* 75 (4):643–669.

Epstein, Larry G. 2008. "Living with Risk." *The Review of Economic Studies* 75 (4):1121–1141.

Epstein, Larry G. and Yoram Halevy. 2020. “Hard-to-Interpret Signals.” Tech. rep., Mimeo, University of Toronto.

Epstein, Larry G and Martin Schneider. 2003. “Recursive multiple-priors.” *Journal of Economic Theory* 113 (1):1–31.

Esponda, Ignacio and Emanuel Vespa. 2021. “Contingent Preferences and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory.” Tech. rep., Mimeo, UCSB.

Freeman, David J, Yoram Halevy, and Terri Kneeland. 2019. “Eliciting Risk Preferences using Choice Lists.” *Quantitative Economics* 10 (1):217–237.

Galanis, Spyros. 2021. “Dynamic Consistency, Valuable Information and Subjective Beliefs.” *Economic Theory* 71 (4):1467–1497.

Ghirardato, Paolo, Fabio Maccheroni, and Massimo Marinacci. 2004. “Differentiating Ambiguity and Ambiguity Attitude.” *Journal of Economic Theory* 118 (2):133–173.

Gilboa, Itzhak and Ehud Lehrer. 1991. “The Value of Information – An Axiomatic Approach.” *Journal of Mathematical Economics* 20 (5):443–459.

Gilboa, Itzhak and David Schmeidler. 1989. “Maxmin Expected Utility with Non-Unique Prior.” *Journal of Mathematical Economics* 18:141–153.

Golman, Russell, David Hagmann, and George Loewenstein. 2017. “Information Avoidance.” *Journal of Economic Literature* 55 (1):96–135.

Grossman, Zachary. 2014. “Strategic Ignorance and the Robustness of Social Preferences.” *Management Science* 60 (11):2659–2665.

Grossman, Zachary and Joël J van der Weele. 2017. “Self-image and Willful Ignorance in Social Decisions.” *Journal of the European Economic Association* 15 (1):173–217.

Gul, Faruk and Wolfgang Pesendorfer. 2021. “Evaluating Ambiguous Random Variables from Choquet to Maxmin Expected Utility.” *Journal of Economic Theory* 192:105129.

Halevy, Yoram. 2007. “Ellsberg Revisited: An Experimental Study.” *Econometrica* 75 (2):503–536.

Hertwig, Ralph and Christoph Engel. 2016. “Homo Ignorans: Deliberately Choosing not to Know.” *Perspectives on Psychological Science* 11 (3):359–372.

Ho, Emily H, David Hagmann, and George Loewenstein. 2021. “Measuring Information Preferences.” *Management Science* 67 (1):126–145.

Holt, Charles A. 1986. “Preference Reversals and the Independence Axiom.” *The American Economic Review* 76 (3):508–515.

Holt, Charles A and Susan K Laury. 2002. “Risk Aversion and Incentive Effects.” *American Economic Review* 92 (5):1644–1655.

Huck, Steffen and Georg Weizsäcker. 1999. “Risk, Complexity, and Deviations from Expected-Value Maximization: Results of a Lottery Choice Experiment.” *Journal of Economic Psychology* 20 (6):699–715.

Ivanenko, Victor and Valery Labkovskii. 1986. “A Class of Criterion-Choosing Rules.” *Dokl. Akad. Nauk SSSR* 287:564–567. [Sov. Phys. Dokl. 31(3) (1986)].

Jonas, Eva, Stefan Schulz-Hardt, Dieter Frey, and Norman Thelen. 2001. “Confirmation Bias in Sequential Information Search After Preliminary Decisions: An Expansion of Dissonance Theoretical Research on Selective Exposure to Information.” *Journal of Personality and Social Psychology* 80 (4):557–571.

Karlsson, Niklas, George Loewenstein, and Duane Seppi. 2009. “The Ostrich Effect: Selective Attention to Information.” *Journal of Risk and Uncertainty* 38:95–115.

Karni, Edi and Zvi Safra. 1987. “”Preference Reversal” and the Observability of Preferences by Experimental Methods.” *Econometrica* 55 (3):675–685.

Kellner, Christian, Mark T. Le Quement, and Gerhard Riener. 2022. “Reacting to Ambiguous Messages: An Experimental Analysis.” *Games and Economic Behavior* 136:360–378.

Kocher, Martin G, Amrei Marie Lahno, and Stefan T Trautmann. 2018. “Ambiguity Aversion is not Universal.” *European Economic Review* 101:268–283.

König-Kersting, Christian, Christopher Kops, and Stefan T. Trautmann. 2023. “A Test of (Weak) Certainty Independence.” *Journal of Economic Theory* 209:105623.

Kosfeld, Michael and Ulrich Schüwer. 2017. “Add-on Pricing in Retail Financial Markets and the Fallacies of Consumer Education.” *Review of Finance* 21 (3):1189–1216.

Kőszegi, Botond. 2003. “Health Anxiety and Patient Behavior.” *Journal of Health Economics* 22 (6):1073–1084.

Krähmer, Daniel and Rebecca Stone. 2013. “Anticipated Regret as an Explanation of Uncertainty Aversion.” *Economic Theory* 52 (2):709–728.

Lacko, James M and Janis K Pappalardo. 2010. “The Failure and Promise of Mandated Consumer Mortgage Disclosures: Evidence from Qualitative Interviews and a Controlled Experiment with Mortgage Borrowers.” *American Economic Review* 100 (2):516–521.

Li, Jian. 2020. “Preferences for Partial Information and Ambiguity.” *Theoretical Economics* 15 (3):1059–1094.

Liang, Yucheng. 2021. "Learning from Unknown Information Sources." Tech. rep., Mimeo, Stanford University.

List, John A. 2007. "On the Interpretation of Giving in Dictator Games." *Journal of Political economy* 115 (3):482–493.

Meechan, Geraldine, John Collins, and Keith Petrie. 2002. "Delay in Seeking Medical Care for Self-Detected Breast Symptoms in New Zealand Women." *The New Zealand Medical Journal (Online)* 115 (1166).

Oechssler, Jörg and Alex Roomets. 2014. "Unintended Hedging in Ambiguity Experiments." *Economics Letters* 122 (2):243–246.

Oechssler, Jörg and Alex Roomets. 2015. "A Test of Mechanical Ambiguity." *Journal of Economic Behavior & Organization* 119:153–162.

Oster, Emily, Ira Shoulson, and E Dorsey. 2013. "Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease." *American Economic Review* 103 (2):804–830.

Pavan, Alessandro and Jean Tirole. 2022. "Expectation Conformity in Strategic Cognition." Tech. rep., Mimeo, Northwestern University.

Riedel, Frank, Jean-Marc Tallon, and Vassili Vergopoulos. 2018. "Dynamically Consistent Preferences under Imprecise Probabilistic Information." *Journal of Mathematical Economics* 79:117–124.

Savage, Leonard J. 1954. *The Foundations of Statistics*. New York: Dover Publications.

Segal, Uzi. 1990. "Two-Stage Lotteries Without the Reduction Axiom." *Econometrica* 58 (2):349–377.

Shishkin, Denis and Pietro Ortoleva. 2023. "Ambiguous Information and Dilation: An Experiment." *Journal of Economic Theory* 208:105610.

Somasundaram, Jeeva and Enrico Diecidue. 2017. "Regret Theory and Risk Attitudes." *Journal of Risk and Uncertainty* 55 (2-3):147–175.

Trautmann, Stefan T and Gijs Van De Kuilen. 2015. "Ambiguity Attitudes." In *Blackwell Handbook of Judgment and Decision Making*, edited by G. Keren and G. Wu. Blackwell, 89–116.

Wakker, Peter. 1988. "Nonexpected Utility as Aversion of Information." *Journal of Behavioral Decision Making* 1 (3):169–175.

Zilker, Veronika, Ralph Hertwig, and Thorsten Pachur. 2020. "Age Differences in Risk Attitude Are Shaped by Option Complexity." *Journal of Experimental Psychology: General* 149 (9):1644–1683.