

Kops, Jan Christopher; Pasichnichenko, Illia

**Working Paper**

## A test of information aversion

Discussion Paper Series, No. 682

**Provided in Cooperation with:**

Alfred Weber Institute, Department of Economics, University of Heidelberg

*Suggested Citation:* Kops, Jan Christopher; Pasichnichenko, Illia (2020) : A test of information aversion, Discussion Paper Series, No. 682, University of Heidelberg, Department of Economics, Heidelberg,  
<https://doi.org/10.11588/heidok.00028244>

This Version is available at:

<https://hdl.handle.net/10419/235005>

**Standard-Nutzungsbedingungen:**

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

**Terms of use:**

*Documents in EconStor may be saved and copied for your personal and scholarly purposes.*

*You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.*

*If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.*

University of Heidelberg

Department of Economics



Discussion Paper Series | No. 682

**A Test of Information Aversion**

---

Christopher Kops and Illia Pasichnichenko

---

April 2020

# A Test of Information Aversion\*

Christopher Kops, Illia Pasichnichenko<sup>†</sup>

## Abstract

The standard Bayesian model implies that information can never have a negative value. We put this implication to the proof. Our paper provides the first test of the value (positive or negative) of information under uncertainty. We show that the “Bayesian implication” stands in conflict with the information-averse behavior that is revealed in our experiment. This behavior demonstrates that the value of truthful and unambiguous information may indeed be negative. Our findings complement predictions from recent theoretical work in showing that negative value of information correlates with ambiguity aversion. This highlights the importance of counseling for decision-making under uncertainty.

**JEL codes:** D81, D83, D90

**Keywords:** Value of Information, Ambiguity Aversion, Ellsberg paradox, Ellsberg urn

## 1 Introduction

Information is a means to resolve uncertainty. It has economic value. Examples run from individual decision-making under uncertainty, to problems of adverse selection and moral hazard in markets with asymmetric information. A lesser known characteristic of information is that it may just as well increase uncertainty. It may even do so to the point that it impedes the decision-making process and that people are inclined to avoid it. This paper proposes a way to test for information aversion

---

\*This draft: April 16, 2020

<sup>†</sup>Kops, and Pasichnichenko: Heidelberg University, Germany (e-mails: christopher.kops@awi.uni-heidelberg.de and illia.pasichnichenko@awi.uni-heidelberg.de). Correspondence address: Department of Economics, University of Heidelberg, Bergheimer Straße 58, 69115 Heidelberg, Germany. We are grateful to Aurélien Baillon, Peter Duersch, Jürgen Eichberger, Hannes Rau, and Christoph Vanberg for helpful comments and suggestions. Illia Pasichnichenko acknowledges financial support of the Alexander von Humboldt Foundation.

under uncertainty and presents the results from an experimental implementation of that test.

To the best of our knowledge, we are the first to put the value of truthful information under uncertainty to a test. Earlier studies have focussed on other aspects of decision-making in dynamic environments. A related and very recent strand of the literature studies the value of ambiguous information (Epstein and Halevy, 2019; Liang, 2019; Kellner, Le Quement, and Gerhard, 2019; Shishkin and Ortoleva, 2019). Another strand of the literature put consequentialism and dynamic consistency to a test under uncertainty (Cohen et al., 2000; Dominiak, Duersch, and Lefort, 2012; Bleichrodt et al., 2019; Esponda and Vespa, 2019). A third strand 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, 2020) and managerial decision-making (Deshpande and Kohli, 1989).

What separates our approach from all these studies is that we provide a test of the value of truthful and unambiguous information under uncertainty. Under the standard Bayesian model, such information can never have a negative value for the decision maker (DM). As the usual argument goes, she can simply ignore it at no cost and make a decision as if it was not available to her. Our results in this paper show that this reasoning stands in conflict with the information-averse behavior revealed by 62% of participants in our experiment. Furthermore, the behavior we observe is neither random, nor resulting from confusion, nor caused by indifference, nor a by-product of dynamic inconsistency, nor is it driven by reasons previously put forward in the literature on information aversion such as anxiety (Kőszegi, 2003; Eliaz and Spiegel, 2006; Epstein, 2008), disappointment aversion (Dillenberger, 2010; Andries and Haddad, 2017), regret aversion (Krähmer and Stone, 2013; Somasundaram and Diecidue, 2017), optimism maintenance (Brunnermeier and Parker, 2005; Oster, Shoulson, and Dorsey, 2013) and belief investments (Jonas et al., 2001; Dana, Weber, and Kuang, 2007). Rather, our experiment shows that information aversion is significantly correlated with ambiguity aversion.

This result confirms the predictions of recent theoretical work (Snow, 2010; Heyen and Wiesenfarth, 2015; Li, 2019; Galanis, 2019) suggesting a close connection be-

tween ambiguity aversion and unwillingness to receive partial information. As such, ambiguity aversion may be a prerequisite for information aversion in the first place. The real-world and the economics literature are replete with examples supporting this conclusion. To substantiate this claim and suggest possible applications of our results, a few examples may be instructive at this point.

**Example 1.1 (Financial Crisis of 2007–2008).** Consider the study on information search by Fischer et al. (2011). In their experiment, participants were given information suggesting an 80% (or, 20%) chance for the repercussions of the Financial Crisis of 2007–2008 to get worse in the near future. Then, participants had to provide their own predictions for it. Before they were asked to do so again, they received expert statements on the topic and were able to choose between statements confirming or contradicting the initial information. Results showed a clear preference for confirmatory information. Our paper suggests a new interpretation for this phenomenon, one that is entirely based on ambiguity aversion. Knowing the chances are 80% (or, 20%), participants avoid contradicting information, precisely because it makes the situation less clear, less predictable, and more ambiguous.

**Example 1.2 (Creation-evolution Controversy).** According to a 2014 Gallup survey, more than four in ten Americans believe that God created humans in their present form. In the scientific community, evolution by natural selection is accepted as fact. As such, the theory of evolution and all evidence supporting it is partial information for an atheist view of the world, at best. Science leaves plenty of room for God to exist. Why is it then that over 40% of Americans avoid the information supporting evolution by natural selection? The results from our experiment suggest that the reason for this may again be ambiguity aversion. People, who are convinced that certain views of the world are true, avoid information debunking part of it as untrue, because this creates uncertainty about the world and what to believe in.

**Example 1.3 (Echo Chambers).** Echo chambers are a metaphorical description of situations in which people only perceive and seek out information which reinforces their existing views. The phenomenon was invoked to explain outcomes of presidential elections in the United States (Barberá et al., 2015) and of the 2016 Brexit referendum in the United Kingdom (Del Vicario et al., 2017). Again, one of the reasons why this phenomenon persists may be because people are averse to the increase in ambiguity that information contradicting their worldview can create.

New information and its (potentially negative) value are also at the heart of the

literature on growing awareness (see, among others, Karni and Vierø (2013), Galanis (2015), Karni and Vierø (2017)). In the traditional Bayesian framework, new information can only shrink the state space and every new observation may rule out links previously thought possible. The literature on growing awareness, on the other hand, also considers information which enlarges the state space and opens up new possibilities never thought of before. Adding to both strands of the literature, in our experiment, we focus on information that clearly shrinks the state space.

As for our choice of a laboratory experiment, note that the primitive concept in economics for modeling individual behavior, beliefs and attitudes is choice. Choices are indeed at the heart of any microeconomic model that aims at delivering quantitative and qualitative predictions about human behavior. Therefore, a clean test of information aversion should be entirely based on choice data. Specifically, data on choices between a situation with information and the same situation but without the information. The issue with this lies in “unknowing” the information from one situation to the next. While this is nearly impossible to guarantee in the real world, for a test in the laboratory, we can set up two identical situations, add information that pertains to one situation, but not to the other, and ask participants to choose their preferred situation. This, we feel, is a strength of our test. It allows participants to experience and familiarize themselves with both situations before they have to make their choice between the two.

Our findings have important implications for public policy. They suggest that the majority of people reject information even when material benefits are attached to it. And, that this continues to be true when they can experience and familiarize themselves with both the “informed” and “uninformed” situation before they have to make their decision between the two. As stated before, our results show that this information-averse behavior is significantly correlated with ambiguity aversion. Thus, in situations where people perceive subjective ambiguity, extensive counseling may be required at a very early stage for people to fully grasp the consequences of their actions. Even earlier and broader than what is common practice. To come full circle to real-world instances of information aversion alluded to before, this finding corroborates, for example, the important role of counseling for genetic tests (Oster, Shoulson, and Dorsey, 2013), career counseling (Deshpande and Kohli, 1989), or, financial literacy (Karlsson, Loewenstein, and Seppi, 2009).

The next section establishes the link between theory and experiment. The following two sections present the experimental design and the results, highlighting the correlation between information aversion and ambiguity aversion. The final two sections discuss our findings in relation to prominent ambiguity models in the decision theory literature and draw implications for future theoretical developments.

## 2 A Direct Test

For our experiment, we implemented a direct test of information aversion under uncertainty that builds on a dynamic version of an Ellsberg-type urn. As stated in the Introduction, our design allows us to directly compare “informed” and “uninformed” decisions. To see this, we first illustrate in this section that Savage acts (Savage, 1954) can be represented in a straightforward manner using bets on Ellsberg-type urns.

Let  $S$  be a finite set of states. Subsets of  $S$  are referred to as events, i.e., any  $E \subseteq S$  is an event. Let  $X$  denote a set of outcomes. An act  $f$  is a function from  $S$  into  $X$ . We consider a decision maker (DM) who has preferences  $\succsim$  over the set of all possible acts  $\mathcal{F}$ . The DM’s preferences conditional on the occurrence of some event  $E \subset S$  are denoted by  $\succsim_E$ . Let  $V : \mathcal{F} \rightarrow \mathbb{R}$  be a representation of the DM’s preferences  $\succsim$ , and, similarly,  $V_E$  be a representation of the DM’s conditional preferences  $\succsim_E$ .

This definition of the standard framework by Savage (1954) has become a workhorse for theoretical models of decision-making under uncertainty. With it and with the representations of the DM’s beliefs and preferences, we can define the concepts of the value of information, dynamic consistency and ambiguity attitude that we test in our experiment.

For our experimental test of these concepts, we explicitly pin down states and outcomes. To this end, let  $S = \{1^G, 2^B, 3^G, 4^B\}$  and  $X = \{\text{€}4, \text{€}4.5, \text{€}10, \text{€}10.5\}$ . From now on,  $\mathcal{F}$  denotes the set of possible acts between these two sets. An example of an act is  $f$  such that  $f(s)$  is equal to  $\text{€}10$  if  $s = 1^G$ , and to  $\text{€}4$  otherwise.


Our experimental design allows for a direct test of the concepts mentioned above

and defined in detail below, by representing acts as bets on urns. Consider the two urns, Urn  $U$  and Urn  $K$ , in Figure 1. The information available about their compositions is identical for both urns. Each 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 (state  $1^G$ ) and five blue balls marked with the number 2 (state  $2^B$ ). Each of the remaining eleven balls is either green and marked with a 3 (state  $3^G$ ), or, it is blue and marked with a 4 (state  $4^B$ ). The exact number of balls marked with the number 3 is unknown, as is the exact number of balls marked with the number 4. However, taken together, there are exactly eleven balls marked with a 3 or a 4 in each urn.

The difference between Urn  $U$  and Urn  $K$  lies in the information about the color of the randomly drawn ball from it. For any bet on Urn  $U$ , the color of the randomly drawn ball is not known upfront. For any bet on Urn  $K$ , on the other hand, the color of the randomly drawn ball is revealed before participants make their choices.

It is straightforward to see that acts in  $\mathcal{F}$  can be represented by a bet on the number of a randomly drawn ball from one of these urns. Our example-act  $f$  above can be represented as a bet on Urn  $U$  that pays €10 if the drawn ball is marked with number 1 (i.e., if state  $1^G$  occurs) and €4 otherwise (i.e., if any of the states  $2^B$ ,  $3^G$ , or,  $4^B$  occurs).

Figure 1: Urns used to represent acts





## 2.1 Information Aversion

In the framework set up so far, future information about the state of the world can be represented as a partition of  $S$ . The future information of the drawn ball being green or blue, for example, may be represented as the partition of  $S$  into the two events  $G = \{1^G, 3^G\}$  and  $B = \{2^B, 4^B\}$ .

Next, consider some act  $f \in \mathcal{F}$ . Then, the value of the future information  $\{G, B\}$  for this act  $f$  is defined as

$$VOI = [p(G)V_G(f) + p(B)V_B(f)] - V(f),$$

where  $p$  is the DM's prior belief about the states of the world  $S$ .

As illustrated before, any act  $f \in \mathcal{F}$  can be represented as a bet on the number of a randomly drawn ball from Urn  $U$ . Offering the same bet on Urn  $K$ , we can represent the act  $f$  conditional on knowing upfront whether the color of the drawn ball is green or blue, i.e., conditional on the event  $G$ , or,  $B$ .

Following the exposition above, our setup allows to elicit whether the DM's value of the future information  $\{G, B\}$  for act  $f$  is negative. If the DM prefers the bet on Urn  $U$  to the bet on Urn  $K$  both conditional on knowing that the color of the randomly drawn ball is green and blue, i.e.  $V_G(f) < V(f)$  and  $V_B(f) < V(f)$ , then the value of the information  $\{G, B\}$  for this bet is negative.

## 2.2 Dynamic Consistency

This concept establishes a link between conditional and unconditional preferences. In a nutshell, it requires that choices made ex-ante are consistently implemented in the future. In particular, taken together the preferences  $f \succ g$ ,  $g \succ_G f$ , and  $g \succ_B f$  are dynamically inconsistent. Clearly, the ex-ante preferences  $f \succ g$  are inconsistent with the two conditional preferences taken together, because the ex-ante preferences are not implemented in any subset of the partition of  $S$  into  $G$  and  $B$ .

As for the test of information aversion, both acts  $f$  and  $g$  can be represented as a bet on the number of a randomly drawn ball from Urn  $U$ . Offering the same bets

on Urn  $K$ , we can represent the act  $f$  (resp.,  $g$ ) conditional on knowing the color of the drawn ball.

Following the exposition above, our setup allows to elicit whether the DM's preferences are dynamically inconsistent. If the DM prefers the bet  $f$  to the bet  $g$  when offered on Urn  $U$ , but prefers  $g$  to  $f$  when offered on Urn  $K$  conditional on knowing that the color of the randomly drawn ball is green and also conditional on knowing that the color of the randomly drawn ball is blue, then the DM's preferences are dynamically inconsistent.

### 2.3 Ambiguity Attitude

Following the ambiguity-literature (for a summary of experimental work on this topic, see Oechssler and Roomets (2015) and Trautmann and Van De Kuilen (2015)), we test for the DM's ambiguity attitude by offering two simple choice problems. Once, the choice is between two low likelihood bets of which one is risky and the other is ambiguous. Once, the choice is between two high likelihood bets of which one is ambiguous and the other is risky. By changing payoffs in one of the states, the low likelihood risky bet in the first choice problem becomes the high likelihood ambiguous bet in the second choice problem and vice versa. As for the different ambiguity attitudes, an ambiguity averse individual, for instance, chooses a risky bet in both choice problems.

### 2.4 Underlying Assumptions

To keep the experiment brief and reduce the number of choice problems that participants in our experiment have to answer as much as possible, we impose the following assumptions on their choice behavior.

**A1. (State Space)** Choices between bets on Urn  $U$  and Urn  $K$  can be described using the same decision-theoretic model with the same set of states  $S = \{1^G, 2^B, 3^G, 4^B\}$ , the same set of outcomes  $X = \{\text{€}4, \text{€}4.5, \text{€}10, \text{€}10.5\}$ , and the same complete and transitive DM's preference relation  $\succsim$  on the set  $\mathcal{F}$  of acts, i.e. functions from  $S$  to  $X$ .

Under this assumption, all participants are indifferent between a bet on Urn  $U$  and the same bet on Urn  $K$  (before the color is revealed), since the two bets are indistinguishable in the model. Any participant, therefore, has the same belief about the composition of each urn. We impose this assumption on our choice data and control for its validity with a questionnaire at the end of our experiment.

**A2. (Color Symmetry)** Choices between acts that are symmetric w.r.t. color do not depend on the color of the ball drawn from Urn  $K$ . An act is symmetric w.r.t. color if it is both constant on  $\{1^G, 2^B\}$  and on  $\{3^G, 4^B\}$ .

This assumption implies that participants' beliefs about the composition of the 11 unknown balls marked with 3 and 4 are symmetric, which mirrors the symmetry of the information available to participants in our experiment.

**A3. (Translation Invariance)** For all acts  $f, g \in \mathcal{F}$ ,  $f \succsim g$  if and only if  $f + \text{€}0.5 \succsim g + \text{€}0.5$ . In other words, increasing payoffs by €0.5 in each state of the world does not affect a participant's choice between the two corresponding bets.

For a relatively low increase in payoffs as is the case in our experiment, preferences presumably satisfy translation invariance. For relatively large payoff increases this assumption may become problematic, as was recently shown by Baillon and Placido (2019) and König-Kersting, Kops, and Trautmann (2020).

### 3 Experimental Design

For our experiment, we implemented a test of information aversion under uncertainty that builds on the association between acts and bets on urns laid out in Section 2. More specifically, our experiment consists of five choice problems. Three of them were designed to test information aversion under uncertainty and dynamic consistency. The other two choice problems elicit participants' ambiguity preferences. In each choice problem, participants are asked to choose between two different urn-bets. Each choice problem specifies the urn the bets pertain to. It lays out what is known about the composition of balls in the urn. And it reveals how much each bet pays depending on the color and number of a randomly drawn

ball from the urn.

All bets in our experiment refer to the two urns, Urn  $U$  and Urn  $K$ , in Figure 1. To reiterate, the difference between Urn  $U$  and Urn  $K$  lies in the information about the color of the randomly drawn ball from it. For any bet on Urn  $U$ , the color of the randomly drawn ball is not known. It is only known that this color is either green or blue. For any bet on Urn  $K$ , on the other hand, the color of the randomly drawn ball is revealed before participants have to make their choices.

### 3.1 Decision Tasks Designed to Test Information Aversion

In Choice Problem 1 participants are asked to choose between the two bets  $f_1$  and  $g_1$  on a randomly drawn ball from Urn  $U$ , i.e. not knowing the color of the drawn ball. Table 1 specifies what each bet pays depending on the number that the randomly drawn ball is marked with.

Table 1: Bets on Urn  $U$  (ball color unknown)

|       | 5 balls |       | 11 balls |       |
|-------|---------|-------|----------|-------|
|       | $1^G$   | $2^B$ | $3^G$    | $4^B$ |
| $f_1$ | €10     | €10   | €4       | €4    |
| $g_1$ | €4      | €4    | €10      | €10   |

Notes:  $1^G$  is the event that the drawn ball is green and marked with the number 1,  $2^B$  the event of a blue drawn ball marked with a 2, etc.

For Choice Problem 2, a ball is randomly drawn from Urn  $K$  before participants are asked to make their decision between the two urn-bets of this choice problem. Participants receive the information about the color of the drawn ball, but the number it is marked with remains unknown to them. Say, the color of the randomly drawn ball is green. Then, participants receive this information and can infer from their information about Urn  $K$  that this ball can only be marked with a 1 or a 3. Table 2 specifies what each bet pays depending on the color and number that the randomly drawn ball is marked with. In our example of a green drawn ball,  $f_2$  pays €10.5 if this ball is marked with a 1 and €4.5 if it is marked with a 3.  $g_2$ , on the other hand, pays €4.5 if this ball is marked with a 1 and €10.5 if it is marked with a 3.

Table 2: Bets on Urn  $K$  (ball color known)

|       | 5 balls |       | 11 balls |       |
|-------|---------|-------|----------|-------|
|       | $1^G$   | $2^B$ | $3^G$    | $4^B$ |
| $f_2$ | €10.5   | €10.5 | €4.5     | €4.5  |
| $g_2$ | €4.5    | €4.5  | €10.5    | €10.5 |

In Choice Problem 3, participants are asked to choose between their answers to the first and the second choice problem. That is, they can choose between the bet on Urn  $U$  they have chosen in the first choice problem and the bet on Urn  $K$  they have chosen in the second choice problem. Let  $c(f_i, g_i)$  specify a participant's bet chosen in choice problem  $i$ , for  $i = 1, 2$ . Then, in the third choice problem, this participant is asked to choose between

$$c(f_1, g_1) \quad \text{vs.} \quad c(f_2, g_2)$$

Say, a participant has chosen  $g_1$  in Choice Problem 1, i.e.,  $c(f_1, g_1) = g_1$ , and  $g_2$  in Choice Problem 2, i.e.,  $c(f_2, g_2) = g_2$ . Then, in the third choice problem, she is asked to choose between  $g_1$  and  $g_2$ . Note that  $g_1$  still pertains to Urn  $U$  where nothing is known about the color of the randomly drawn ball. On the other hand,  $g_2$  still pertains to Urn  $K$  where the color of the randomly drawn ball is known upfront.

For our interpretation of participants' choices in this decision task, it is important that the first two assumptions of Section 2.4 hold. As mentioned before, the question regarding participants' beliefs about urn-compositions in our questionnaire serves as a control for whether these assumptions about participants' beliefs are justified.

We can classify participants choosing bet  $c(f_1, g_1)$  (resp., bet  $c(f_2, g_2)$ ) as showing a negative value of information (resp., a non-negative value of information). Note that  $f_2 = f_1 + \text{€}0.5$  and  $g_2 = g_1 + \text{€}0.5$ . Therefore, a preference for  $c(f_1, g_1)$  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 in Choice Problem 1 and 2 are symmetric w.r.t. color, this observation does not depend on whether the color of the ball that was actually drawn is green or blue.


It is straightforward to see that the information about the color of the randomly drawn ball reduces the (objective) uncertainty about the true state of the world. Not knowing the ball color, the true state of the world lies in the set  $\{1^G, 2^B, 3^G, 4^B\}$ . Knowing the color of the randomly drawn ball, say it is green, the true state of the world lies in the set  $\{1^G, 3^G\}$ . A clear reduction in uncertainty. On the other hand, given how the bets in Choice Problem 1 align with the underlying information about the number of balls marked with numbers 3 or 4, information about the color of the randomly drawn ball may increase the (subjective) uncertainty attached to these bets.

To see this, observe that, for the first choice problem, while there is uncertainty about the exact number of balls marked with a 3 (resp., with a 4), the winning probabilities under each bet in this choice problem are objectively given. The probability of winning €10 is  $\frac{10}{21}$  under act  $f_1$  and  $\frac{11}{21}$  under act  $g_1$ . For the second choice problem, this is not true. With the information about the color of the randomly drawn ball, winning probabilities under each bet in this choice problem are not objectively given anymore. What is objectively known is that the probability of winning €10.5 lies in the interval  $[\frac{5}{16}, 1]$  under act  $f_2$  and in the interval  $[0, \frac{11}{16}]$  under act  $g_2$ . Hence, information about the ball color may clearly lead to an increase in subjective uncertainty and, therefore, push an individual who is averse to such an increase to pass on the benefits that the second choice problem involves.

Figure 2 shows the negative value of the information  $\{G, B\}$  for an ambiguity averse DM in Choice Problem 1. The DM is assumed to have maxmin expected utility preferences (MEU) and to be risk neutral (see Section 5.2 for the calculations). Note that a DM who is considerably ambiguity averse is willing to forego the €0.5 bonus payment offered in Choice Problem 2 in order to keep ambiguity to a minimum by avoiding the ex ante information about the ball color.

Finally, consider the following normatively appealing, Bayesian-like reasoning: First, the draw of a, say green, ball from Urn  $K$ , suggests that this draw was more likely than that of a blue ball from this urn. This is even true when the assumption of color symmetry does not hold and beliefs about color are not too asymmetric. A participant in this case should choose  $g_2$  over  $f_2$ , because the green draw suggests larger winning probabilities for  $g_2$  than for  $f_2$ . Note that this is even true in the case of a blue draw. Next, the participant should also choose  $g_2$  over  $c(f_1, g_1)$ ,

Figure 2: Value of the information  $\{G, B\}$  in MEU



because payoffs are larger and the draw of a green ball suggest that the winning probabilities for  $g_2$  are not lower than the ones for  $c(f_1, g_1)$ .

### 3.2 Decision Tasks Designed to Test Dynamic Consistency

Under our assumptions from Section 2.4, we can classify participants choosing bet  $f_1$  from Choice Problem 1 and bet  $f_2$  from Choice Problem 2, or, bets  $g_1$  and  $g_2$  as dynamically consistent. All other participants, i.e., those choosing  $f_1$  and  $g_2$ , or,  $g_1$  and  $f_2$ , we can classify as dynamically inconsistent.

To see why preferences for  $f_1$  and  $g_2$  are dynamically inconsistent, note that  $f_2 = f_1 + \text{€}0.5$  and  $g_2 = g_1 + \text{€}0.5$ . Hence, by translation invariance,  $f_1 \succ g_1$  implies  $f_2 \succ g_2$ . If, say, a green ball is drawn from Urn  $K$  for Choice Problem 2, by color symmetry,  $g_2 \succ_G f_2$  implies  $g_2 \succ_B f_2$ . Taken together, the preferences  $f_2 \succ g_2$ ,  $g_2 \succ_G f_2$ , and  $g_2 \succ_B f_2$  are dynamically inconsistent. The same reasoning also establishes that the preferences for  $g_1$  and  $f_2$  are dynamically inconsistent, as well.

### 3.3 Decision Tasks Designed to Test Ambiguity Aversion

In Choice Problem 4 and 5, participants choose between two bets on Urn  $U$  (see Table 3). Note that the bets in these two choice problems differ only in their payoffs in state  $4^B$ . By the usual interpretation, participants choosing bets  $f_4$  and  $g_5$  (resp.,

$g_4$  and  $f_5$ ) are classified as ambiguity averse (resp., seeking). All other participants, i.e., those choosing  $f_4$  and  $f_5$ , or,  $g_4$  and  $g_5$ , are classified as ambiguity neutral.

Table 3: Bets on Urn  $U$  (ball color unknown)

|       | 5 balls |       | 11 balls |       |
|-------|---------|-------|----------|-------|
|       | $1^G$   | $2^B$ | $3^G$    | $4^B$ |
| $f_4$ | €10     | €4    | €4       | €4    |
| $g_4$ | €4      | €4    | €10      | €4    |
| $f_5$ | €10     | €4    | €4       | €10   |
| $g_5$ | €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 of ambiguous balls ( $3^G$  and  $4^B$  balls) is one larger than the number of risky balls ( $1^G$  and  $2^B$ ). This extra ball acts as a tie breaker. Hence, preference for  $f_4$  over  $g_4$  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 4.5, we use this measure as a robustness check in the analysis of our results.

### 3.4 Implementation & Lab Procedures

The experiment was conducted in November and December 2019 in the AWI Lab at the University of Heidelberg. 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 up to €10.5 from the decision tasks. The experiment took about 45 minutes, for which participants earned, on average, €12.8. Before each session, the boxes were checked to contain the correct distribution of colored, marked balls.<sup>1</sup>

<sup>1</sup>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



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 from Urn  $K$  was resolved by drawing the ball physically with the help of a randomly selected participant. Its 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 uncertainty about which of the five choice problems determined participants' payoffs was resolved with the help of another randomly selected participant. Final payoffs were calculated, participants were paid and dismissed from the lab.<sup>2</sup>

### 3.5 Summary Statistics

In total, 115 subjects participated in the experiment. Participants' average age was 23.2 years, and the share of economics students was 25.2%. With 56.5%, the share of female participants is reasonably close to 50% in our experiment.

## 4 Results

### 4.1 Negative VOI

Our main research question is whether subjects do assign a negative value to information under uncertainty. Such participants would choose their solution to Choice Problem 1 over their solution to Choice Problem 2, when it comes to their choice between these two bets in Choice Problem 3. Table 4 shows the percentage of participants that made negative VOI-choices. The main result of our study is very

---

four. Furthermore, instead of urns we used cardboxes.

<sup>2</sup>All files necessary for replicating the experiment and the results are available on University of Heidelberg's data repository at <https://doi.org/10.11588/data/GORNAZ>. Instructions were translated from German. Original instructions are available from the authors upon request.

clear: at 61.7%, the share of participants showing a negative value of information is substantial.

Table 4: Share of Negative VOI

| Choice Problem 3                | Share           | N   |
|---------------------------------|-----------------|-----|
| $c(f_1, g_1)$ vs. $c(f_2, g_2)$ | 61.7% VOIneg ** | 115 |

Notes: VOIneg = negative value of information. Column “Share” shows percentage of  $c(f_1, g_1)$  chosen in Choice Problem 3. Binary choice: Two-sided binomial test against  $p = 0.5$ . \*\* denotes significance at 5%.

As discussed in Section 3, our conclusion above rests on participants having similar beliefs about the compositions of both urns. In a questionnaire at the end of the experiment, we asked participants for their estimates of the number of green balls in each urn. For those participants who showed negative VOI and for those who did not, we calculated the average distance between estimates of the number of green balls in Urn  $U$  and in Urn  $K$ . In line with our assumption of similar beliefs about urn-compositions, there was no significant difference between the estimates of the number of green balls for each urn ( $p = 0.208$ , two-sided t-test).

Motivated by recent theoretical work (Li, 2019; Eichberger and Pasichnichenko, 2020), we checked whether negative VOI is correlated with ambiguity aversion. Indeed, Table 5 shows that subjects making negative VOI-choices were significantly more (often) ambiguity averse than subjects making non-negative VOI-choices. On the other hand, the latter group was significantly more (often) ambiguity neutral than the group making negative VOI-choices. Among ambiguity averse subjects, 76.9% showed a negative value of information. Thus, our findings are in line with the predictions of recent theoretical contributions. The fact that we were able to replicate the relationship between VOI and ambiguity attitude suggested by (Li, 2019) counteracts worries that our results stem from subjects’ choices being random or driven by confusion. This lends further credence to the robustness of our results.

We asked participants about their confidence in their choices for each decision task. Participants who made negative VOI-choices were significantly less confident in their choices for Choice Problem 2 than in their choices for Choice Problem 1 compared to participants who made non-negative VOI-choices (difference in confidence between decision tasks: VOIneg 1.14 vs. VOInonneg 0.41,  $p = 0.002$ , two-sided

Table 5: Negative VOI, Ambiguity Attitude and Dynamic Inconsistency

|                       | All   | Non-negative VOI | Negative VOI |          |
|-----------------------|-------|------------------|--------------|----------|
|                       | Share | Share            | Share        | p-values |
| Ambiguity averse      | 33.9% | 20.5%            | 42.3%        | 0.012 ** |
| Ambiguity neutral     | 59.1% | 70.5%            | 52.1%        | 0.048 ** |
| Ambiguity seeking     | 7.0%  | 9.0%             | 5.6%         | 0.506    |
| Dynamic inconsistency | 35.7% | 29.5%            | 39.4%        | 0.279    |

Notes:  $N = 115$ . Non-negative VOI: 44; negative VOI: 71. Columns “Share” show percentages in these two subgroups of ambiguity averse ( $f_4 \succ g_4$  and  $g_5 \succ f_5$ ), neutral ( $f_4 \succ g_4$  and  $f_5 \succ g_5$ , or,  $g_4 \succ f_4$  and  $g_5 \succ f_5$ ) and seeking choices ( $g_4 \succ f_4$  and  $f_5 \succ g_5$ ), as well as dynamically inconsistent choices ( $f_1 \succ g_1$  and  $g_2 \succ f_2$ , or,  $g_1 \succ f_1$  and  $f_2 \succ g_2$ ). Two-sided t-test. \*\* denotes significance at 5%.

t-test). The larger drop in confidence for neg-VOI participants is exactly what we would expect from participants who prefer to avoid partial information.

Finally, in our data, there is no significant correlation between negative value of information and dynamic inconsistency. Table 5 shows that the percentage of dynamically inconsistent choices is larger among participants who made negative VOI-choices than among those who made non-negative VOI-choices. But this difference is not statistically significant.

## 4.2 Ambiguity Attitudes

We find some support for the common finding of a fourfold pattern of ambiguity attitudes (Kocher, Lahno, and Trautmann, 2018) restricted to the gain domain. That is, we observe statistically significant ambiguity aversion for high likelihood gain prospects. At the same time, we observe that choices imply attitudes closer to ambiguity seeking or neutrality for the case of low likelihood gains (high 82.6% (95) vs. low 44.3% (51),  $p < 0.001$ , two-sided t-test). The share of risky choices is significantly lower for low likelihood gains than for high likelihood gains. The fact that we replicate large parts of the fourfold pattern of ambiguity attitudes restricted to the gain domain lends further credence to the robustness of our results.

Table 6: Gain-domain Part of Fourfold Pattern

|                      | High (52%) |        | Low (24%) |         |
|----------------------|------------|--------|-----------|---------|
| <i>Gain</i>          |            |        |           |         |
| Choice Problem 4 & 5 | 82.6%      | AA *** | 44.3%     | AS n.s. |

Notes: High / Low designate the probabilities of gain realizations. The cells show percentages of risky prospects chosen and the ambiguity attitude implied: AA = ambiguity averse; AS = ambiguity seeking. Binary choice: Two-sided binomial test against  $p = 0.5$ . \*\*\* denotes significance at 1%.

### 4.3 Dynamic Consistency

According to participants’ responses to Choice Problems 1 and 2, we can classify them as dynamically consistent or not. Table 7 shows that 64.3% of all participants in our experiment are dynamically consistent. This result changes when we look at the subgroup of participants who are ambiguity averse according to the last two decision tasks. In line with the results by Dominiak, Duersch, and Lefort (2012), ambiguity averse participants are significantly more (often) dynamically inconsistent (51.3% vs. 27.6%,  $p = 0.016$ , two-sided t-test). On the other hand, ambiguity neutral participants are significantly more (often) dynamically consistent (75.0% vs. 48.9%,  $p = 0.005$ , two-sided t-test).

Table 7: Share of Dynamically Consistent Choices

| Decision Tasks 3 & 4                | Share |        | N   |
|-------------------------------------|-------|--------|-----|
| $f_1 \succ g_1$ iff $f_2 \succ g_2$ | 64.3% | DC *** | 115 |

Notes: DC = dynamic consistency. Column “Share” shows percentage of choices  $f_1$  and  $f_2$  (resp.,  $g_1$  and  $g_2$ ) in Choice Problems 1 and 2. Binary choice: Two-sided binomial test against  $p = 0.5$ . \*\*\* denotes significance at 1%.

### 4.4 Gender Effect

Finally, we also find a gender effect insofar as the share of males is significantly smaller among those participants who are information-averse than among those who are not (male: VOIneg 35.2% vs. VOInonneg 54.6%,  $p = 0.045$ ; two-sided

t-test).

## 4.5 Indifference

As was already mentioned, we did not offer an indifferent option. However, additional to each decision, participants were asked to mark “How strong is your preference for the alternative you choose?” on a scale ranging from 1 (very weak) to 5 (very strong). Subjects who marked one could be interpreted as having no confidence that their choices are better than the alternatives, that is, as being indifferent. As a robustness check, we discard them from the analysis. All results stay valid with two exceptions. That is, 1) Table 5, ambiguity neutral: VOIneg vs. VOInonneg,  $p = 0.058$ , i.e. lower level of significance, 2) Section 4.3, dynamic inconsistency: ambiguity averse vs. not ambiguity averse,  $p = 0.067$ , i.e. lower level of significance.

## 5 Theoretical Discussion

### 5.1 Subjective Expected Utility

As stated in the Introduction, information can never have a negative value for a Bayesian DM, since it is assumed that she maximizes subjective expected utility. In particular, consider the value of information  $\{G, B\}$  in Choice Problem 1:

$$\begin{aligned}
 VOI = p(G) \max_{s \in G} \frac{p(s)}{p(G)} u(h(s)) + p(B) \max_{s \in B} \frac{p(s)}{p(B)} u(h(s)) \\
 - \max_{s \in S} p(s) u(h(s)),
 \end{aligned}$$

where the maximums are taken over  $h \in \{f_1, g_1\}$ . One can see that the value of information is non-negative for any prior belief  $p$ .

In particular, assume  $p(1^G) = \frac{5}{21}$ ,  $p(2^B) = \frac{5}{21}$ , and the uniform prior on  $\{3^G, 4^B\}$  under the principle of insufficient reason, i.e.,  $p(3^G) = p(4^B) = \frac{5.5}{21}$ . Then all three maxima are attained at  $g_1$ . Clearly, the information  $\{G, B\}$  is of zero value to a

Bayesian DM. Hence, she would never forego the increase in payments offered in Choice Problem 2.

## 5.2 Maxmin Expected Utility

Non-Bayesian DMs may be averse to information (Wakker, 1988). Note that while there is only one way to be Bayesian, there are many ways to be non-Bayesian. This section invokes maxmin expected utility (Ivanenko and Labkovsky, 1986; Gilboa and Schmeidler, 1989) as the alternative approach and discusses the results from our experiment within the MEU-framework. More specifically, we consider an ambiguity averse DM with MEU-preferences. Such a DM's beliefs form a set of prior probability distributions on the state space and she ranks acts by maximizing the minimal expected utility with respect to this set of priors.

To apply MEU to our experiment, let the DM's set of priors over  $S = \{1^G, 2^B, 3^G, 4^B\}$  be given by

$$\mathcal{C} = \left\{ (p(1^G), p(2^B), p(3^G), p(4^B)) = \left( \frac{5}{21}, \frac{5}{21}, \frac{5.5+x}{21}, \frac{5.5-x}{21} \right) : -\varepsilon \leq x \leq \varepsilon \right\}$$

Here,  $0 \leq \varepsilon \leq 5.5$  measures DM's subjective ambiguity. In other words,  $\varepsilon$  captures the range of the number of  $3^G$ -balls that she deems possible. For instance, if  $\varepsilon = 1.5$ , then the DM thinks that the urn contains between four and seven  $3^G$ -balls.

To simplify the exposition, we assume that  $u(\text{€}4) = 0$  and  $u(\text{€}10) = 1$ . Then, for the maxmin expected utility,  $I(\cdot)$ , of acts  $f_4$  and  $g_4$  from Table 3, we have

$$I(f_4) = \min_{p \in \mathcal{C}} \sum_{s \in S} p(s)u(f_4(s)) = \frac{5}{21}$$

$$I(g_4) = \min_{p \in \mathcal{C}} \sum_{s \in S} p(s)u(g_4(s)) = \min_{-\varepsilon \leq x \leq \varepsilon} \frac{5.5+x}{21} = \frac{5.5-\varepsilon}{21}$$

Thus,  $f_4 \succ g_4$  if  $\varepsilon > 0.5$ . On the other hand, since

$$I(f_5) = \min_{-\varepsilon \leq x \leq \varepsilon} \frac{5 + 5.5 - x}{21} = \frac{10.5 - \varepsilon}{21}$$

$$I(g_5) = \min_{-\varepsilon \leq x \leq \varepsilon} \frac{5.5 + x + 5.5 - x}{21} = \frac{11}{21}$$

this implies that  $g_5 \succ f_5$ . In this way, MEU can explain ambiguity averse choices in Choice Problem 4 and 5 of our experiment.

In Choice Problem 1, the payoffs under each act nicely align with the uncertainty such that  $I(f_1) = \frac{10}{21}$  and  $I(g_1) = \frac{11}{21}$ , and so  $g_1 \succ f_1$ . Suppose now that while choosing between  $f_1$  and  $g_1$ , the DM is offered information, whether the drawn ball is green or blue. To estimate the value of such information, we first calculate the values of the two hypothetical choice problems – the choice between  $f_1$  and  $g_1$  conditional on the event  $G$  and the choice between  $f_1$  and  $g_1$  conditional on the event  $B$ . To do this, we need to update the DM's priors and calculate the minimal expected utilities of the two acts. The Full Bayesian update of  $\mathcal{C}$  to the event  $G$  is given by

$$\mathcal{C}_G = \left\{ (p(1^G), p(2^B), p(3^G), p(4^B)) = \left( \frac{5}{10.5+x}, 0, \frac{5.5+x}{10.5+x}, 0 \right) : -\varepsilon \leq x \leq \varepsilon \right\}$$

Then, for the conditional maxmin expected utility,  $I_G$ , of  $f_1$  and  $g_1$ , we have

$$I_G(f_1) = \min_{p \in \mathcal{C}_G} \sum_{s \in S} p(s)u(f_1(s)) = \min_{-\varepsilon \leq x \leq \varepsilon} \frac{5}{10.5+x} = \frac{5}{10.5+\varepsilon}$$

$$I_G(g_1) = \min_{p \in \mathcal{C}_G} \sum_{s \in S} p(s)u(g_1(s)) = \min_{-\varepsilon \leq x \leq \varepsilon} \frac{5.5+x}{10.5+x} = \frac{5.5-\varepsilon}{10.5-\varepsilon}$$

Therefore,  $f_1 \succ_G g_1$  if  $\varepsilon \geq \varepsilon^* = \frac{\sqrt{21}}{2} \approx 2.3$  and  $g_1 \succ_G f_1$  otherwise. The value of the choice problem,  $V_G$ , is the utility of the best act, i.e.,

$$V_G = \max \{I_G(f_1), I_G(g_1)\} = \begin{cases} I_G(f_1), & \text{if } \varepsilon \geq \varepsilon^* \\ I_G(g_1), & \text{otherwise} \end{cases}$$

By symmetry,  $V_B = V_G$ .

Finally, the value of information about the color of the drawn ball is equal to the minimal expected increase in the value of the choice problem with respect to the set of priors  $\mathcal{C}$ , i.e.,

$$\text{VOI} = \min_{p \in \mathcal{C}} [p(G)V_G + p(B)V_B] - V_0.$$

Since  $V_G = V_B$ , we have

$$\text{VOI} = V_G - V_0.$$

Note that the value of the original choice problem is given by

$$V_0 = \max\{I(f_1), I(g_1)\} = \frac{11}{21}$$

If  $\varepsilon^* \leq \varepsilon \leq 5.5$ , then

$$\text{VOI} = I_G(f_1) - V_0 = \frac{10}{21 + 2\varepsilon} - \frac{11}{21} < 0$$

On the other hand, if  $0 \leq \varepsilon < \varepsilon^*$ , then

$$\text{VOI} = I_G(g_1) - V_0 = \frac{11 - 2\varepsilon}{21 - 2\varepsilon} - \frac{11}{21} < 0$$

Therefore, an ambiguity averse DM with MEU-preferences is averse to information about the color of the drawn ball. Only in the extreme case (of a Bayesian) where  $\mathcal{C}$  is a singleton set, we have  $\varepsilon = 0$  and  $\text{VOI} = 0$ .

Recall that the parameter  $\varepsilon$  measures a DM's subjective ambiguity. The more ambiguity averse a DM is, the more money she will pay/forego to avoid the information. For any  $\varepsilon$ , we can calculate the value of information  $\{G, B\}$  and, by assuming risk neutrality, convert utility units to monetary amounts. Figure 2 in Section 3.1 graphically illustrates this relationship. Put differently, it illustrates the monetary amount required to compensate an ambiguity averse DM with MEU-preferences for the negative value of the information  $\{G, B\}$ . For example, if the DM thinks that the urn contains between four and seven  $3^G$ -balls, which corresponds to  $\varepsilon = 1.5$ , then she will choose a bet on urn  $K$  only if the bonus payment is at least €0.48.

## 6 Conclusion

We set out to study the value of information. Specifically, the value of partial information under uncertainty. We provide the first test of whether this value can be negative or not. The findings from our experiment show that at 61.7% the percentage of individuals showing a negative value of information is very substantial and robust. This is not to say that the value of information is always negative. Rather, it shows that if information is partial and has the potential to increase subjective uncertainty, then people will be inclined to “pay” for not having the information. Figure 2 shows there is an upper limit as to what price people are willing to pay for such ignorance.



Furthermore, our results show that information-averse individuals are significantly more ambiguity averse. The economics and psychology literature has previously put forward reasons for information aversion such as disappointment aversion, regret aversion, optimism maintenance and belief investments (Golman, Hagmann, and Loewenstein, 2017). Our results cannot confirm that these are the reasons behind the information aversion we observe. Rather, our experiment shows that information aversion is significantly correlated with ambiguity aversion.

Finally, our theoretical and experimental results have important implications for public policy. In particular, they suggest that many people reject information even when monetary benefits are attached to it. And, that this continues to be true when they can experience and familiarize themselves with both the “informed” and “uninformed” situation. Furthermore, our results show that this information-averse behavior is significantly correlated with ambiguity aversion. Thus, in situations where people perceive subjective ambiguity, extensive counseling may be required at a very early stage for people to fully grasp the consequences of their actions. To come full circle to real-world instances of information aversion alluded to in the Introduction, this finding corroborates, for example, the important role of counseling for genetic tests (Oster, Shoulson, and Dorsey, 2013), financial literacy (Karlsson, Loewenstein, and Seppi, 2009), or, career counseling (Deshpande and Kohli, 1989).

# Appendix A

Table A1: Correlations

|        | 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

- Andries, Marianne and Valentin Haddad. 2017. “Information Aversion.” Tech. rep., NBER Working Paper No. w23958.
- Baillon, Aurélien and Lætitia Placido. 2019. “Testing Constant Absolute and Relative Ambiguity Aversion.” *Journal of Economic Theory* 181:309–332.
- Barberá, Pablo, John T Jost, Jonathan Nagler, Joshua A Tucker, and Richard Bonneau. 2015. “Tweeting from Left to Right: Is Online Political Communication more than an Echo Chamber?” *Psychological Science* 26 (10):1531–1542.
- Bleichrodt, Han, Jürgen Eichberger, Simon Grant, David Kelsey, and Chen Li. 2019. “A Test of Dynamic Consistency and Consequentialism in the Presence of Ambiguity.” Tech. rep., Discussion Papers 1803, University of Exeter, Department of Economics.
- 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.
- Brunnermeier, Markus K and Jonathan A Parker. 2005. “Optimal Expectations.” *American Economic Review* 95 (4):1092–1118.
- 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.
- Del Vicario, Michela, Fabiana Zollo, Guido Caldarelli, Antonio Scala, and Walter Quattrociocchi. 2017. “Mapping Social Dynamics on Facebook: The Brexit Debate.” *Social Networks* 50:6–16.
- 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.
- Eichberger, Jürgen and Illia Pasichnichenko. 2020. "Decision-Making with Partial Information." Discussion Paper Series 680, University of Heidelberg, Alfred Weber Institut.
- Eliasz, Kfir and Ran Spiegler. 2006. "Can Anticipatory Feelings Explain Anomalous Choices of Information Sources?" *Games and Economic Behavior* 56 (1):87–104.
- Epstein, Larry G. 2008. "Living with Risk." *The Review of Economic Studies* 75 (4):1121–1141.
- Epstein, Larry G. and Yoram Halevy. 2019. "Hard-to-Interpret Signals." Tech. rep., Mimeo, University of Toronto.
- Esponda, Ignacio and Emanuel Vespa. 2019. "Contingent Preferences and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory." Tech. rep., Mimeo, UCSB.
- Fischer, Peter, Andreas Kastenmüller, Tobias Greitemeyer, Julia Fischer, Dieter Frey, and David Crelley. 2011. "Threat and Selective Exposure: The Moderating Role of Threat and Decision Context on Confirmatory Information Search after Decisions." *Journal of Experimental Psychology: General* 140 (1):51–62.
- Galanis, Spyros. 2015. "The Value of Information under Unawareness." *Journal of Economic Theory* 157:384–396.
- . 2019. "Dynamic Consistency, Valuable Information and Subjective Beliefs." Tech. rep., Discussion Papers, University of London.
- Gilboa, Itzhak and David Schmeidler. 1989. "Maxmin Expected Utility with Non-Unique Prior." *Journal of Mathematical Economics* 18:141–153.
- Golman, Russell, David Haggmann, and George Loewenstein. 2017. "Information Avoidance." *Journal of Economic Literature* 55 (1):96–135.
- Hertwig, Ralph and Christoph Engel. 2016. "Homo Ignorans: Deliberately Choosing not to Know." *Perspectives on Psychological Science* 11 (3):359–372.

- Heyen, Daniel and Boris R Wiesenfarth. 2015. “Informativeness of Experiments for MEU - A Recursive Definition.” *Journal of Mathematical Economics* 57:28–30.
- Ho, Emily H, David Hagmann, and George Loewenstein. 2020. “Measuring Information Preferences.” *Management Science* :forthcoming.
- Ivanenko, Victor and Valery Labkovsky. 1986. “On the Functional Dependence between the Available Information and the Chosen Optimality Principle.” In *Stochastic Optimization*. Springer, 388–392.
- 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 (2):95–115.
- Karni, Edi and Marie-Louise Vierø. 2013. ““Reverse Bayesianism”: A Choice-based Theory of Growing Awareness.” *American Economic Review* 103:2790–2810.
- . 2017. “Awareness of Unawareness: A Theory of Decision Making in the Face of Ignorance.” *Journal of Economic Theory* 168:301–328.
- Kellner, Christian, Mark T. Le Quement, and Riener Gerhard. 2019. “Reacting to Ambiguous Messages: An Experimental Analysis.” Tech. rep., Mimeo, University of Southampton.
- 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. 2020. “A Test of (Weak) Certainty Independence.” Tech. rep., Mimeo, University Heidelberg.
- 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.

- Li, Jian. 2019. “Preferences for Partial Information and Ambiguity.” Tech. rep., Discussion Papers, McGill University.
- Liang, Yucheng. 2019. “Learning from Unknown Information Sources.” Tech. rep., Mimeo, Stanford University.
- 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.
- Savage, Leonard J. 1954. *The Foundations of Statistics*. New York: Dover Publications.
- Shishkin, Denis and Pietro Ortoleva. 2019. “Ambiguous Information and Dilation: An Experiment.” Tech. rep., Mimeo, Princeton University.
- Snow, Arthur. 2010. “Ambiguity and the Value of Information.” *Journal of Risk and Uncertainty* 40 (2):133–145.
- 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.