Available online at www.sciencedirect.com

ScienceDirect

Journal of Economic Theory 208 (2023) 105610

JOURNAL OF Economic Theory

www.elsevier.com/locate/jet

Ambiguous information and dilation: An experiment ☆

Denis Shishkin [a,*] ⓡ Pietro Ortoleva [b]

[a] *Department of Economics, University of California San Diego, United States of America*
[b] *Department of Economics and SPIA, Princeton University, United States of America*

Received 8 October 2021; final version received 10 December 2022; accepted 14 January 2023
Available online 21 January 2023

Abstract

With standard models of updating under ambiguity, new information may *increase* the amount of relevant ambiguity: the set of beliefs may 'dilate.' We experimentally test one sharp case: agents bet on a risky urn and get information that is truthful or not based on the draw from an Ellsberg urn. With common models, the set of beliefs dilates, and the value of bets decreases for ambiguity-averse agents and increases for ambiguity-seeking ones. Instead, we find that the value of bets does not change for ambiguity-averse individuals, while it increases substantially for ambiguity-seeking ones. We also test bets on ambiguous urns, in which case we find sizable reactions to ambiguous information.

JEL classification: C91; D81; D90

Keywords: Updating; Ambiguous information; Ambiguity aversion; Ellsberg paradox; Maxmin expected utility

☆ The order of authors was randomized; the authors contributed equally. We thank the editor Pierpaolo Battigalli, an associate editor, and two anonymous referees for the very valuable comments and advice, as well as Marina Agranov, Simone Cerreia-Vioglio, Mark Dean, David Dillenberger, Adam Dominiak, Yoram Halevy, Faruk Gul, Yucheng Liang, Fabio Maccheroni, Massimo Marinacci, Stephen Morris, Wolfgang Pesendorfer, Leeat Yariv, Sevgi Yuksel, and various audiences for useful discussions and suggestions. We are grateful to Rui Tang for outstanding research assistance. Ortoleva gratefully acknowledges the financial support of NSF Grant SES-1763326.

* Corresponding author.
E-mail addresses: dshishkin@ucsd.edu (D. Shishkin), pietro.ortoleva@princeton.edu (P. Ortoleva).

https://doi.org/10.1016/j.jet.2023.105610

1. Introduction

Extensive theoretical and experimental literature has studied modeling and implications of ambiguity—when payoffs depend on states of the world for which there is no objective probability distribution. One important aspect is how ambiguity interacts with updating. However, the theoretical literature in this area has not reached a consensus, and experimental analysis is much more limited.

We experimentally study one aspect of ambiguity and updating: how agents react to *information of ambiguous reliability*. Agents make bets and receive messages that can be truthful or misleading depending on the draw from an Ellsberg urn. We study this scenario for three reasons.

First, this experiment tests a key implication of standard models of updating under ambiguity, namely that information may *increase* relevant ambiguity and make agents *worse off*—the so-called 'dilation' of sets of beliefs. For example, with widespread updating rules for MaxMin Expected Utility, the set of relevant beliefs may become larger (dilate) after information. This makes ambiguity-averse agents strictly *worse off* after information, and they should be willing to *pay to avoid it*. Appealing or not, this is an implication of widespread models that, to our knowledge, has not been tested. Our experiment with information of ambiguous reliability allows us to test a very sharp case, in which dilation should occur after any message—'all news is bad news' (Gul and Pesendorfer, 2021).

Second, the dilation feature of updating that we test is central to many applications to game theory and mechanism design. In fact, the very type of ambiguous information that we study in our experiment is used in applications. Below, we discuss papers on ambiguity in mechanism design and Bayesian persuasion in which, in the motivating examples of both, the ambiguity pertains to whether messages are truthful or not, exactly as in our experiment.

Third, our experiment contributes to the growing interest in ambiguous information, focusing on the case of ambiguity in the reliability of messages. While ambiguity in informativeness may be commonplace in real life, it is studied only by few recent papers as discussed below.

Experiment Subjects first evaluate bets on the color of a ball drawn from an urn. In some cases, they receive a message about the winning color; this message is truthful or misleading depending on the draw from a (separate) 2-color Ellsberg bag of chips. After subjects acknowledge the message, we measure how the value of bets changes. We also measure the (positive or negative) value of this information. In some questions, the payoff-relevant draw is made from a risky, 50/50 urn; thus, all ambiguity is in the message. In other questions, draws are made from an ambiguous urn. We also measure subjects' ambiguity-aversion.

Relation to theories Standard models of updating under ambiguity make clear predictions. Consider the MaxMin Expected Utility model (MMEU) of Gilboa and Schmeidler (1989). Two updating rules are widespread: Full-Bayesian (FB), wherein the posterior set includes updates of all beliefs; and Maximum-Likelihood (ML), wherein the updated set includes only beliefs that satisfy a maximum-likelihood criterion. When the payoff-relevant state is risky, the relevant set of beliefs without information is a singleton. But, after ambiguous information, as in our experiment, if subjects are not ambiguity neutral with FB or ML, the relevant set of beliefs *dilates* and is no longer a singleton. Because of the ambiguity in information, bets on the risky urn *become ambiguous*. Ambiguity-averse agents should then decrease the value of bets after information and even pay to avoid it. (The opposite is true for ambiguity-seeking agents.) Note that these predictions hold for *any* message, meaning that all news is bad news.

One may well consider some of these predictions unappealing—especially that information must make ambiguity-averse individuals strictly worse off. In particular, this view is in line with the critique suggested in the statistics literature (Good, 1967, 1974; Seidenfeld, 1981; Walley, 1991). However, these implications are central predictions of widely used theories and play an important role in application.

We also show that these predictions hold beyond MMEU: they also apply to Bayesian updating of (symmetric) Smooth Ambiguity model (Klibanoff et al., 2005, 2009). The fact that this model vastly differs from MMEU points to the pervasiveness of this implication.

Results When the payoff-relevant state is risky, ambiguity-averse or neutral subjects typically do *not* change the value of bets after information. The median change is zero, and the majority has exactly zero change. For ambiguity-averse subjects, there is also no robust relation between ambiguity-aversion and the size of change (or the probability it is non-zero). They also typically give zero value to information. All these findings contrast to the theoretical predictions of *negative* reaction to information and negative value of information.

On the other hand, ambiguity-seeking subjects typically increase their valuation after information, and this change in value is strongly related to their ambiguity affinity, in line with the theoretical prediction. Yet, many still value the information close to zero.

When the payoff-relevant state is ambiguous, ambiguity-averse subjects slightly increase valuations, while ambiguity-seeking subjects decrease them. Theories make no prediction for this case.

Implications We found that ambiguity-averse agents do *not* react negatively to our type of ambiguous information, in contrast to the dilation property of widespread models of updating under ambiguity. We thus conclude the paper with a discussion of alternative rules that are instead compatible with our findings.

First, subjects may be using FB or ML, but complement it by strategically choosing *whether* to process the information. While under Subjective Expected Utility information is always weakly valuable—there is no benefit in ignoring it—this is no longer the case under ambiguity. Thus, subjects may decide to ignore information when it is harmful. To our knowledge, the only existing model with this feature is Dynamically Consistent Updating (Hanany and Klibanoff, 2007, 2009), wherein agents form an ex-ante optimal plan contingent on information and use an updating rule such that they want to implement it. This rule, however, violates consequentialism (unrealized parts of the decision problem may influence beliefs).

Alternatively, subjects may follow Proxy Updating (Gul and Pesendorfer, 2021) or Contraction Updating (Tang, 2022), which satisfy consequentialism but restrict dilation. Proxy Updating is designed precisely to rule out instances of 'all news is bad news.' While the exact model cannot be applied to our setup—it is defined for totally monotone capacities, which is not the case here—our results are aligned with the underlying idea of this approach. Contraction Updating is defined for a general MMEU, and it is compatible with our main finding that ambiguity-averse agents do not change their valuation of bets after information.

Literature A theoretical literature discusses updating rules under ambiguity (Gilboa and Marinacci, 2013, Sec. 5), while an experimental one tested dynamic consistency and consequentialism (Cohen et al., 2000; Dominiak et al., 2012; Bleichrodt et al., 2021; Esponda and Vespa, 2021), how sampling from ambiguous sources affects ambiguity preferences (Ert and Trautmann, 2014), learning from sequences of observations (Moreno and Rosokha, 2016), in groups (De Filippis et al., 2022), or from stock prices (Baillon et al., 2017).

Several recent papers also study ambiguous information (Epstein and Halevy, 2022; Liang, 2022; Kellner et al., 2022; Kops and Pasichnichenko, 2022).[1] A common difference with our work is that they do not study the dilation property, our primary focus. Building on Epstein and Schneider (2007, 2008), Epstein and Halevy (2022) define and characterize attitudes to 'signal ambiguity,' the ambiguity on the informativeness of a signal, and tests it experimentally.[2] Focusing on a setup with ambiguity, they find that signal ambiguity significantly increases deviations from Bayesian updating. While related to our work in its interest in ambiguous signals, in Epstein and Halevy (2022) the payoff-relevant state is ambiguous and signals are always informative, but the agent does not know how much. On the other hand, in our experiment, the payoff-relevant state can be risky and the ambiguity is on *whether* the signal is informative or misleading. The papers are thus complementary: our design is less extensive on ambiguous information but allows us to test the dilation property and a form of ambiguous information used in the applied literature.

A contemporaneous paper by Liang (2022) studies updating with risky state under simple and uncertain (ambiguous and compound) signals, as well as ambiguous and compound state under simple signals. It compares updating under different types of signals that correspond to the same average simple signal and finds that subjects under-react to uncertain information, which is more pronounced for good news rather than bad news. Also contemporaneous, Kellner et al. (2022) study messages with ambiguous reliability but asymmetric and with three messages, one of which is informative. They find a relation between reactions and ambiguity attitude and a similar reaction to ambiguous and compound-risk signals. Their design does not allow for tests of dilation. Kops and Pasichnichenko (2022) tests negative value of ambiguous information and finds that many subjects are information averse in the sense that their value of information in a binary choice task is at most -0.5 euros. Similar to our results, the fraction of information-averse subjects is significant across ambiguity attitudes, and no strong relation exists between negative value of information and ambiguity attitude (more than half of information-averse subjects are not ambiguity-averse). However, they report a much larger fraction of information-averse subjects, a dissimilarity that may be due to experimental design differences.

Lastly, as noted above, the ambiguous information we study is used in applications of models of ambiguity to strategic environments. Bose and Renou (2014) study mechanism design wherein the allocation stage is preceded by an ambiguous mediated communication stage. Beauchêne et al. (2019) study Bayesian persuasion with an ambiguity-averse receiver and a sender who can commit to ambiguous signals. Both papers assume FB, and some of their results are linked to its dilation property. Pahlke (2022) and Cheng (2021) study ambiguous persuasion under alternative updating rules that account for dynamic consistency and dilation.

[1] Moreover, Vinogradov and Makhlouf (2020) augment the Ellsberg experiment with vague statements about an ambiguous payoff-relevant state, which may be perceived as ambiguous signals.

[2] We learned about Epstein and Halevy (2022) before finalizing our design. We thank Yoram Halevy for useful discussions.

2. Theories of updating and ambiguous information

Consider a set of prizes $X = \mathbb{R}$ and a state space $S = \Omega \times M$, where $\Omega = \{R, B\}$ are the payoff-relevant states (colors of the ball), and $M = \{r, b\}$ the messages.[3]

Assume preferences are represented by MMEU: given strictly increasing, continuous utility $u: X \to \mathbb{R}$ and set of beliefs $\Pi \subseteq \Delta(S)$, agents evaluate act $f: S \to X$ by $\min_{\pi \in \Pi} \mathbb{E}_\pi[u \circ f]$ if ambiguity-averse; if ambiguity-seeking, max replaces min.

Given Π, let Π^Ω denote the set of marginals over Ω and identify any $\hat{\pi} \in \Pi^\Omega$ with $\hat{\pi}(R) \in [0, 1]$. In line with our experiment, assume that all beliefs are symmetric in the sense that the likelihood that a message is truthful or not is independent of the payoff-relevant state, i.e., $\pi(r|R) = \pi(b|B)$. Moreover, assume that Π contains the uninformative belief $\bar{\pi}$ (i.e., $\bar{\pi}(r|R) = \bar{\pi}(b|R) = 0.5$) and, if $|\Pi| \neq 1$, it does so in its relative interior.[4] We call *coherent* any closed set of beliefs that satisfies these restrictions.

Updating rules For any event $D \subseteq S$, let Π_D denote the set of beliefs after information D. The following two updating rules are the most common.

Full Bayesian (FB) updating (Wasserman and Kadane, 1990; Jaffray, 1992; Pires, 2002; Ghirardato et al., 2008), the most common in applications, is defined by

$$^{\mathrm{FB}}\Pi_D := \{\pi(\cdot|D) : \pi \in \Pi\}.$$

Maximum Likelihood (ML) updating (Dempster, 1967; Shafer, 1976; Gilboa and Schmeidler, 1993) is defined by

$$^{\mathrm{ML}}\Pi_D := \{\pi(\cdot|D) : \pi \in \underset{\pi' \in \Pi}{\operatorname{argmax}}\, \pi'(D)\}.$$

Under FB, individuals update all beliefs following Bayes' rule. Under ML, they retain (and update) only the beliefs with the highest likelihood of the realized event.

Variables of interest Suppose an individual is ambiguity-averse and considers a bet that pays a baseline of x and adds y if a color chosen by the individual realizes. Its certainty equivalent is

$$c_m(x, y) := u^{-1}\left(\max_{\omega \in \{R,B\}} \min_{\pi \in \Pi_m^\Omega} \pi(\omega) u(x + y) + (1 - \pi(\omega)) u(x) \right),$$

where $m \in \{\varnothing, r, b\}$ denotes either the message received or no information (letting $\Pi_\varnothing := \Pi$).

Our first variable of interest is the *Information Premium*: the difference between certainty equivalents of such bets before and after a message, defined as

[3] The choice of the state space may have some implications in our results below. We follow standard practice and adopt the smallest state space in which payoff-relevant acts are defined (thus, we need the color of the ball extracted for the urn) such that the information can be encoded as a subset of the state space. Smaller state spaces are not possible; we are not aware of how considering larger state spaces could change our conclusions.

[4] Asymmetries in the set of beliefs about the color of the ball extracted would increase the certainty equivalent of all bets (since subjects can choose the color to bet on). If the set of beliefs about information does not include the uninformative belief, the value of the information would necessarily increase: if all beliefs are such that the information is more likely to be truthful, the agent should follow it. However, if all are such that it is more likely to be untruthful, the agent may simply follow the opposite. In either case, the information becomes strictly valuable, increasing the value of all bets after information.

$$P_m := c_m(x, y) - c_{\varnothing}(x, y).$$

Our second variable of interest is the *Value of Information*: the amount V that makes the individual indifferent between no information and receiving a message while modifying payoffs by V. It is defined as the solution of the following equation

$$u(c_{\varnothing}(x, y)) = \min_{\pi \in \Pi} \mathbb{E}_\pi \left[u(c_m(x - V, y))\right].$$

For ambiguity-seeking individuals, all variables are defined analogously, replacing min with max.

Risky payoff-state Suppose that the payoff-relevant state is risky. Then, we can derive the following result (all proofs appear in Appendix B).

Proposition 1. *Consider an MMEU agent with a coherent set of beliefs Π such that the payoff-state is risky ($|\Pi^\Omega| = 1$) and symmetric ($\pi(R) = 0.5, \forall \pi \in \Pi$).*[5] *Then with FB and ML for any $m \in M$:*

1. *if* ambiguity-averse *and* $|\Pi| > 1$*:* $P_m < 0$*,* $V < 0$*;*
2. *if* ambiguity-seeking *and* $|\Pi| > 1$*:* $P_m > 0$*,* $V > 0$*;*
3. *if* ambiguity-neutral *(i.e.,* $|\Pi| = 1$*):* $P_m = V = 0$.

The proposition shows that ambiguity-averse individuals *must* have negative Information Premia P_m for any message m, as well as a negative Value of Information V. Both are positive if ambiguity-seeking.

For intuition, suppose $\Pi = \text{co}(\{\pi_1, \pi_2, \pi_3\})$ with $\pi_1(r|R) = 0.8$ (messages point in the right direction), $\pi_2(r|R) = 0.5$ (messages uninformative), and $\pi_3(r|R) = 0.2$ (messages misleading). Before information, the set of marginals over Ω was a singleton—we have a risky state. But, after information, this set becomes full-dimensional: ${}^{\text{FB}}\Pi_r^\Omega = {}^{\text{FB}}\Pi_b^\Omega = [0.2, 0.8] \supset \{0.5\} = \Pi^\Omega$. The multiplicity of beliefs about the truthfulness of the message generates multiple beliefs about payoff-relevant states. The set of posteriors not only includes the original belief, but also does so in its interior. Importantly, this holds for any message. Seidenfeld and Wasserman (1993) call this property of FB *dilation*. Proposition 1 shows that dilation occurs whenever $|\Pi| > 1$.

With FB, because the set of beliefs dilates, ambiguity-averse individuals have strictly lower certainty equivalents. Moreover, because this holds for any message, the value of information is negative. The opposite holds for ambiguity-seeking individuals.

With ML, individuals focus only on beliefs that maximize the likelihood of the message. When $\pi(R) = 0.5$ for all $\pi \in \Pi$, however, the likelihood of both messages is 0.5. Thus, all beliefs in Π are considered, and ML coincides with FB.

Ambiguous payoff-state Consider now ambiguous payoff-relevant states.

Proposition 2. *Consider an MMEU individual with a coherent set of beliefs Π. Then with FB and ML:*

[5] Symmetry of the payoff-relevant marginal is assumed only for ease of exposition. Without it, predictions with FB are the same; with ML, it remains true that both P and V are not zero unless there is ambiguity-neutrality.

1. *If* $|\Pi| > 1$*, then both* P_m *and* V *can be zero, negative, or positive under both ambiguity-aversion and seeking;*
2. *If* $|\Pi| = 1$*, then* $P_m = V = 0$*,* $\forall m \in M$*.*

When the payoff-relevant state is ambiguous, there are no predictions. Results depend on the set of beliefs across Ω and informativeness. Appendix B.1 contains examples showing dilation and contraction.

Beyond MMEU? Similar results hold beyond MMEU. Consider another popular model, the Smooth Model (Klibanoff et al., 2005). Preferences are represented by

$$U(f) = \mathbb{E}_\mu[\phi(\mathbb{E}_\pi[u \circ f])],$$

where the continuous, strictly increasing $\phi : \mathbb{R} \to \mathbb{R}$ captures ambiguity attitude and $\mu \in \Delta(\Delta(S))$ denotes the belief over beliefs. Following the literature, assume that an individual updates μ following Bayes' rule.[6] Define P_m and V analogously using this model.

Proposition 3. *Consider an individual whose preferences follow the Smooth Model with* μ *symmetric*[7] *and such that* $\text{supp}(\mu)$ *is coherent. If the payoff-state is risky* ($|\text{supp}(\mu)^\Omega| = 1$)*, then for any* $m \in M$*:*

1. *if* ambiguity-averse *(strictly concave* ϕ*):* $P_m < 0$, $V < 0$*;*
2. *if* ambiguity-seeking *(strictly convex* ϕ*):* $P_m > 0$, $V > 0$*;*
3. *if* ambiguity-neutral *(affine* ϕ*):* $P_m = V = 0$.

3. Experiment

3.1. Design

The experiment includes two parts, for a total of 6 questions. In each, subjects were asked to compare fixed amounts of money with a bet on their chosen color drawn from an urn. With two exceptions mentioned below, all bets paid \$20 if the ball was of the chosen color, zero otherwise. Subjects were asked to compare each bet with a list of amounts of money ranging from \$0 to \$20, in a Multiple Price List (MPL). To simplify the task, subjects had to click only once in each list, indicating the point at which to switch from the bet to the amount of money.[8]

Different questions involved urns of two types. Risky urns had a known composition: 100 balls, 50 of each color. Ambiguous urns had 100 balls of two colors with unknown compositions.

Questions were of three kinds. For each, subjects answered one question wherein the payoff-relevant urn was risky, and one in which it was ambiguous.

[6] Identical results also hold if individuals simply update the beliefs in the support (and not the belief over beliefs).

[7] We call belief μ symmetric if for any $Q \subset \Delta(S)$, $\mu(Q) = \mu(0.5 - Q)$, where $0.5 - Q := \{\pi' \in \Delta(S) : \exists \pi \in Q, \forall s \in S, \pi'(s) = 0.5 - \pi(s)\}$.

[8] By monotonicity, subjects should prefer bets against low amounts and 'switch' as the amount grows. The software (oTree; Chen et al., 2016) asked to indicate the point at which to switch. Subjects were also allowed to indicate no switch (always bet or always money). This procedure simplified choice but forced monotonicity. Subjects received extensive instruction and training.

1. Basic Questions. Q1 and Q4 asked subjects to pick a color to bet on and then the certainty equivalent of a \$20 bet using the MPL procedure. In Q1, the urn was risky. In Q4, it was ambiguous. Comparing the answers, we obtained a measure of ambiguity aversion.

2. Information Questions. Q2 and Q5 measured the certainty equivalent of a bet, but after information. At the beginning of the question, the computer drew a ball from the payoff-relevant urn—determining the color that pays the bet—and a chip from a bag with 100 chips of 2 colors and unknown composition. The computer then displayed a message for the subject indicating the color of the ball drawn from the urn. Whether this message was truthful or misleading, however, depended on the chip drawn. If the chip was of one color, the computer told the truth; otherwise, it reported the opposite. In these questions, subjects were first shown the urn, then shown how the message was determined, and finally given the message. They then had to acknowledge it by clicking on the corresponding color. With the message remaining on screen, they had to pick a color to bet on and evaluate the bet using an MPL. In Q2, the payoff-relevant urn was risky; in Q5, it was ambiguous.

3. Information-Value Questions. Q3 and Q6 were similar to the questions above, but also measured the value of information. In these questions, subjects first faced an MPL in which they chose between no information and information, in addition an increase or decrease of their potential winning for the question (from a base of \$20), ranging from −\$5 to \$5. After their respective choice, the computer randomly picked a line from this MPL and implemented their selection. If in that line, the subject chose no information, they proceeded with the evaluation of the bet without it. If they chose the information and a change in payoffs, they received both before evaluating the bet. In Q3, the underlying urn was risky; in Q6, it was ambiguous.

All questions used different urns and different colors, reducing the possibility of hedging across questions. This was clearly explained. Similarly, the bags that determined the information were all different and involved different colors. For symmetry, all colors for urns and bags were randomly selected.[9]

Order and incentives The 6 questions were grouped into two parts. Part I included the 3 questions involving bets on risky urns, in the following order: Q1, the evaluation of a bet on a risky urn; Q2, the evaluation of a bet on a risky urn after information; Q3, the evaluation of a bet on a risky urn after deciding whether or not to receive the information. Part II was identical, but with ambiguous urns. Questions are summarized in Table 1. There were two possible orders: in Order A, Part I then Part II; in Order B, the opposite.

Subjects received a participation fee of \$10 and a completion fee of \$15. One of the 6 questions was randomly selected for payment, and one of the lines of the MPL with the comparison between bets and amounts of money were randomly selected. Subjects received their choice for that line.[10]

9 Colors were selected randomly for each subject and each question from the same set, except that, for each subject, we avoided repetitions and pairings of similar colors.

10 Paying one randomly selected question is incentive compatible under Expected Utility but not beyond; no general incentive compatible mechanism exists (Karni and Safra, 1987; Bade, 2015; Azrieli et al., 2018). Some studies indicate that this may not be a concern (Beattie and Loomes, 1997; Cubitt et al., 1998; Hey and Lee, 2005; Kurata et al., 2009), while others suggest caution (Freeman et al., 2019).

Table 1
Questions.

		Payoff Urn	*Info*
Part I	Q1	Risky	No
	Q2	Risky	Yes
	Q3	Risky	Evaluate
Part II	Q4	Ambiguous	No
	Q5	Ambiguous	Yes
	Q6	Ambiguous	Evaluate

3.2. *Predictions and construction of variables*

We now map the theoretical predictions from Section 2 to our experiment. From the MPLs comparing bets and amounts of money, we approximate the certainty equivalent of each bet. From the MPLs comparing information vs. no information (in Q3 and Q6), we approximate the value of information.

Because, in our experiment, choices involve different urns, we have to make two assumptions. First, that Π before information is the same in questions of the same type, which is justified by the use of identical urns with randomly drawn colors. Second, we assume that subjects' ambiguity attitude is the same across questions and with respect to information- and payoff-relevant states, as we have done implicitly in Section 2. In particular, we assume that, if Π^{Ω} is not a singleton when the payoff-state is ambiguous, then the set of beliefs about the truthfulness of messages is also not a singleton.[11]

We identify ambiguity attitudes by comparing the answer to Q1 (risky urn, no info) and Q4 (ambiguous urn, no info): a higher/equal/lower certainty equivalent in Q1 than in Q4 indicates ambiguity-aversion/neutrality/seeking. The *Ambiguity Premium* is the difference between the value in Q1 and the value in Q4.

The *Information Premium* P is defined as the value in Q2 (Q5) minus the value in Q1 (Q4) for risky (ambiguous) payoff-states. The *Value of Information* V is elicited directly in the first part of Q3 (Q6) for risky (ambiguous) payoff-states.[12]

Applying the results from Section 2, the predictions for risky payoff-states are that both the Information Premium and the Value of Information should be negative/zero/positive for individuals that are ambiguity-averse/neutral/seeking. For ambiguous payoff-states, the only prediction is that both measures should be zero for ambiguity-neutral.

Constructing variables Because MPLs have finite grids, our value elicitation is approximate. Following standard practice, we define the value as the mid-point between the grid points where

[11] This assumption is further justified by the fact that the description of the urn for the ambiguity in the information structure is identical to the description of the ambiguity in the payoff-relevant state, except for bag/jar terminology and randomized colors. For the former, from Q2: "*. . . 100 chips, either Lime or Brown. The composition of the bag is unknown: there may be no Lime chips or no Brown chips, or any other composition.*" For the latter, from Q4 "*. . . 100 Olive and Rose balls. The composition of the jar is unknown: there may be no Olive chips or no Rose chips, or any other composition.*"

[12] Note that the second part of Q3 (Q6) was equivalent to Q1 (Q4) or to Q2 (Q5), depending on whether a subject received information. In such cases, the majority chose consistently.

the switch occurred.[13] However, the true certainty equivalent may be anywhere in that range, and the approximation may matter in computing if variables are equal, smaller, or bigger than zero. We take the following conservative approach. Recall that the Information Premium is the difference between two certainty equivalents, each obtained via an MPL. When computing whether it is above, at, or below zero, we report it in two ways. First, using the procedure above, we denote results by >0, $=0$, <0. Second, we report the percentage of answers that are compatible with zero value, and denote them by ≈ 0.[14]

For the Value of Information, the grid is \$0.1 around 0, and 0 is an option on the grid. By our mid-point construction, no subject can have a value of 0: even if they give 0 value to information, they must have either 0.05 and -0.05. In calculations, we use these numbers; in reporting >0, $=0$, or <0, we put 0.05 and -0.05 in the $=0$ category. Thus, zero values may be overestimated.

Finally, we take a (standard) conservative approach to compute ambiguity attitudes. Namely, we classify as ambiguity-averse or seeking only subjects whose behavior is not compatible with ambiguity-neutrality; thus, subjects who switch in two adjacent lines in Q1 and Q4 are classified as ambiguity-neutral.[15] This implies that we may be overestimating ambiguity-neutral individuals. (As will be clear below, our main conclusions would not change with different classifications.)

3.3. *Results*

91 volunteer undergraduate students participated in 4 sessions of approximately 30 minutes at the PeXL laboratory at Princeton University in February 2019. Average earnings were \$35.2. We eliminated from our analysis 2 subjects who reported strictly dominated answers in multiple questions. Including them does not change any of our conclusions (see Appendix C.2). Recall that we used two different orders. While this had some effect, the patterns we describe hold throughout only with minor differences (Appendix C.1).

The distribution of ambiguity-averse, neutral, and seeking is 35 (39.3%), 37 (41.6%), and 17 (19.1%).[16] Median ambiguity premia are relatively large for both averse (\$2.5) and seeking subjects ($-$\$2). (Table 4 in Appendix C.1 contains all details.)

3.3.1. *Risky payoff-state*

We begin with the case in which payments depend on a draw from a risky urn. Results appear on the left of Table 2 and in Fig. 1. There, the top panel includes a scatter plot of the Information Premium and the Ambiguity Premium. Colors represent ambiguity attitude: red for averse, blue for neutral, green for seeking.[17] On the right is a stacked bar plot depicting the proportions of values that are >0, <0 and $=0$. The bottom panel repeats this for the Value of Information.

Considering all subjects, the mean Information Premium P is positive, but the median is zero. The mean Value of Information V is slightly negative, while the median is compatible with

[13] For example, if the individual chose the bet against \$10 but the next grid point, say \$10.2, against the bet, we set the certainty equivalent at \$10.1.

[14] For example, suppose in Q1 the switch is between \$10 and \$10.2; in Q2, it is one line below, between \$10.2 and \$10.5. With our procedure, values are 10.1 for Q1 and 10.35 for Q2, indicating a positive difference. But this behavior is also compatible with an individual who has zero difference: the true certainty equivalent may be \$10.2 in both questions, but the individual may break indifference in different ways. This behavior is thus marked >0 but also ≈ 0.

[15] Like above, individuals may have the same certainty equivalent but break the indifference in opposite ways.

[16] The fraction of ambiguity-averse is lower than general population results (Chapman et al., forthcoming), but in line with selective universities. Recall that our procedure may overestimate ambiguity-neutral individuals.

[17] For the color version of the figure, the reader is referred to the web version of this article.

Table 2
Results.

ambiguity attitude	Risky payoff state: All	averse	neutral	seeking	Ambiguous payoff state: All	averse	neutral	seeking
				Information Premium P				
theory prediction		<0	$=0$	>0		-	0	-
median	0	0	0	1	0	0	0	−1
mean	0.46	−0.56	0.6	2.2	0.13	0.57	0.08	−0.68
≈ 0	65%	69%	78%	29%	58%	54%	76%	29%
$=0$	56%	57%	68%	29%	52%	49%	70%	18%
>0	28%	11%	30%	59%	24%	34%	14%	24%
<0	16%	31%	3%	12%	25%	17%	16%	59%
				Value of Information V				
theory prediction		<0	$=0$	>0		-	0	-
median	−0.05	−0.05	−0.05	0.05	−0.05	−0.05	−0.05	−0.05
mean	−0.41	−0.23	−0.73	−0.09	−0.43	−0.11	−0.64	−0.61
$=0$	54%	51%	51%	65%	57%	57%	60%	53%
>0	10%	11%	5%	18%	10%	14%	5%	12%
<0	36%	37%	43%	18%	33%	29%	35%	35%
# of obs.	89	35	37	17	89	35	37	17

† Due to rounding, the percentages for >0, <0, $=0$ need not always add up to 100%.

zero. To test the theoretical predictions, however, we have to separate our analysis by ambiguity attitude.

Ambiguity-averse subjects For ambiguity-averse subjects, the median Information Premium is *zero*. The majority (57%) has a value of exactly zero, while 69% have values compatible with zero (denoted ≈ 0). Only 31% have negative values.

A coherent picture emerges with the Value of Information: it is zero for most ambiguity-averse subjects. (Recall that -0.05 is compatible with indifference with zero.) Of the minority with non-zero values, the larger group (37%) has negative values.

We can also test the relation between the degree of ambiguity-aversion and the reaction to information: Are subjects with higher ambiguity premium more likely to have negative, or smaller, information premium? In particular, if the perception of ambiguity about the informational ambiguity is similar to that about the ambiguity in the payoff-relevant state, then FB and ML predict that the higher the subject's ambiguity premium, the higher her reaction to information.

The plots in Fig. 1 suggest that, aside from the outliers in the bottom right, for ambiguity-averse subjects there is no relation between ambiguity and information premia. This is confirmed by statistical analysis. First, for ambiguity-averse subjects, the likelihood of having a non-zero Information Premium is not related to the Ambiguity Premium (Probit, $z=0.69$, $p=0.491$). Second, we can test if Information Premium and Ambiguity Premium are negatively related for ambiguity-averse subjects. Note that, all else equal, our design is potentially biased to generate this (negative) correlation spuriously. Because both variables are constructed using the answer to Q1—the certainty equivalent of the bet on a risky urn without information—noise in this measure would generate a negative spurious relation. (See Section 3.5 for more.)

Despite this fact, in our data, Information Premium and Ambiguity Premium are not robustly related. While an OLS regression does give a relation ($t=-3.65$, $p=0.001$), this is driven by the

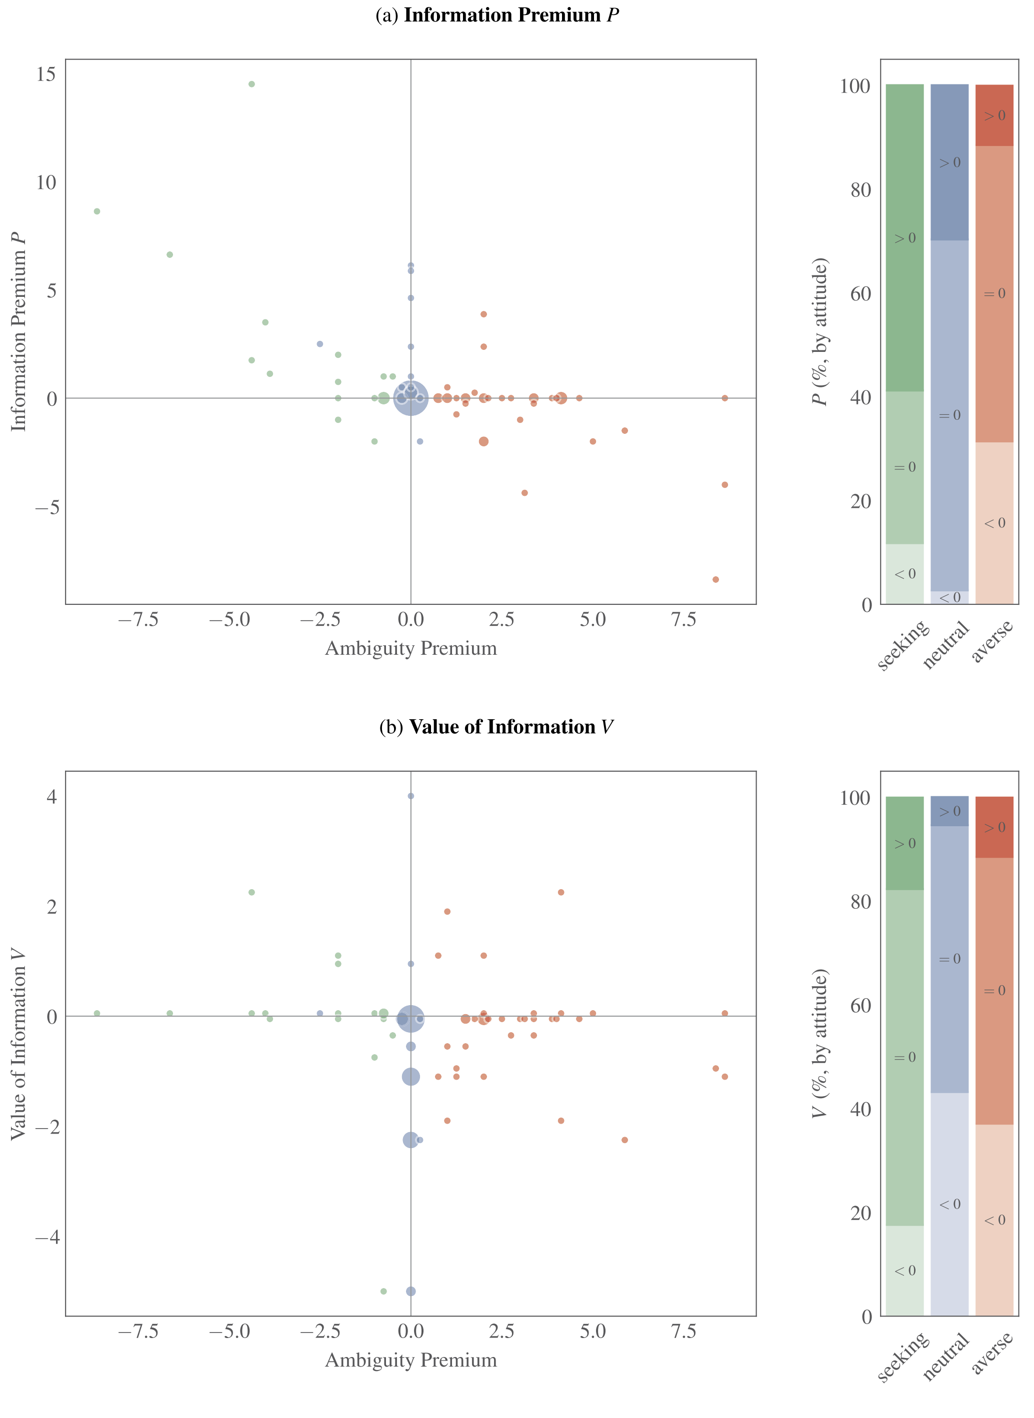

Fig. 1. Results, risky payoff-relevant state, graphically.

two outliers in the bottom right.[18] Eliminating the outliers eliminates the relation ($t = -0.99$, $p = 0.331$). A Quantile regression with all subjects finds no relation ($t = 0.00$, $p = 1$).

Moreover, there is also no relation between the Value of Information and the Ambiguity Premium ($t = 0.91$, $p = 0.369$).

Ambiguity-seeking subjects Patterns are different for ambiguity-seeking subjects: 59% have *positive* Information Premium. Both median and mean are also remarkably high. That is, ambiguity-seeking subjects substantially increase their valuation after ambiguous information.

The Value of Information, however, remains zero for the majority. This suggests that they do not expect to react positively to messages; however, the actually do once confronted with them. While our design does not offer a conclusive explanation for this finding, there exist several plausible reasons for it. For instance, planning how one reacts to information arguably requires more sophistication than simply reacting to it, which could drive the observed discrepancy between V and P. Furthermore, various manifestations of dynamic inconsistencies are commonly observed in experimental data in the context of ambiguity. Recall, however, that our procedure may overestimate how many subjects give zero value to information (Section 3.2).

Fig. 1 also suggests a relation between Ambiguity Premium and Information Premium for ambiguity-seeking subjects. Regressing the two, we find a significant, positive relationship ($t = -3.91$, $p = 0.001$). Note, however, that this relation could be spuriously strengthened by our design, as discussed above. There is no relation with the Value of Information ($t = -1.09$, $p = 0.292$).

Ambiguity-neutral subjects Ambiguity-neutral subjects exhibit a large majority of zero values (non-zeros tend to positive and small). About half have zero Value of Information. Interestingly, 43% give strictly *negative* values, hinting at a non-instrumental role of information.

3.3.2. *Ambiguous payoff-state*

Results for ambiguous payoff-states appear on the right part of Table 2 and in Fig. 2. Clear, but different, patterns emerge.

Considering all subjects, the median Information Premium is again zero, but the mean is slightly positive. The majority still reports zero. Similar results hold for the Value of Information, albeit with a small, negative mean.

Ambiguity-averse subjects have again a median Information Premium close to zero; a sizable fraction has Value of Information either zero or compatible with it. But these are smaller fractions than above, and 34% have strictly *positive* Information Premium.

The opposite pattern holds for ambiguity-seeking subjects: now the majority (59%) has *negative* Information Premium. The Value of Information, however, remains predominantly zero in both cases. Ambiguity-neutral subjects, unsurprisingly, exhibit patterns similar to those found with risky urns.

Overall, we have a positive relationship between the Information Premium and the Ambiguity Premium ($t = 2.73$, $p = 0.008$), but this does not hold separately for ambiguity-averse ($t = 0.95$, $p = 0.349$) or seeking ($t = 1.80$, $p = 0.091$) subjects. There is also no relation with the Value of Information ($t = 0.97$ overall; $t = 1.44$ for averse; $t = -1.15$ for seeking).

[18] This is due to their extreme response in Q1, where both give 18.5 as the certainty-equivalent of a 50/50 bet \$20/\$0, giving large values for both measures and generating the relation.

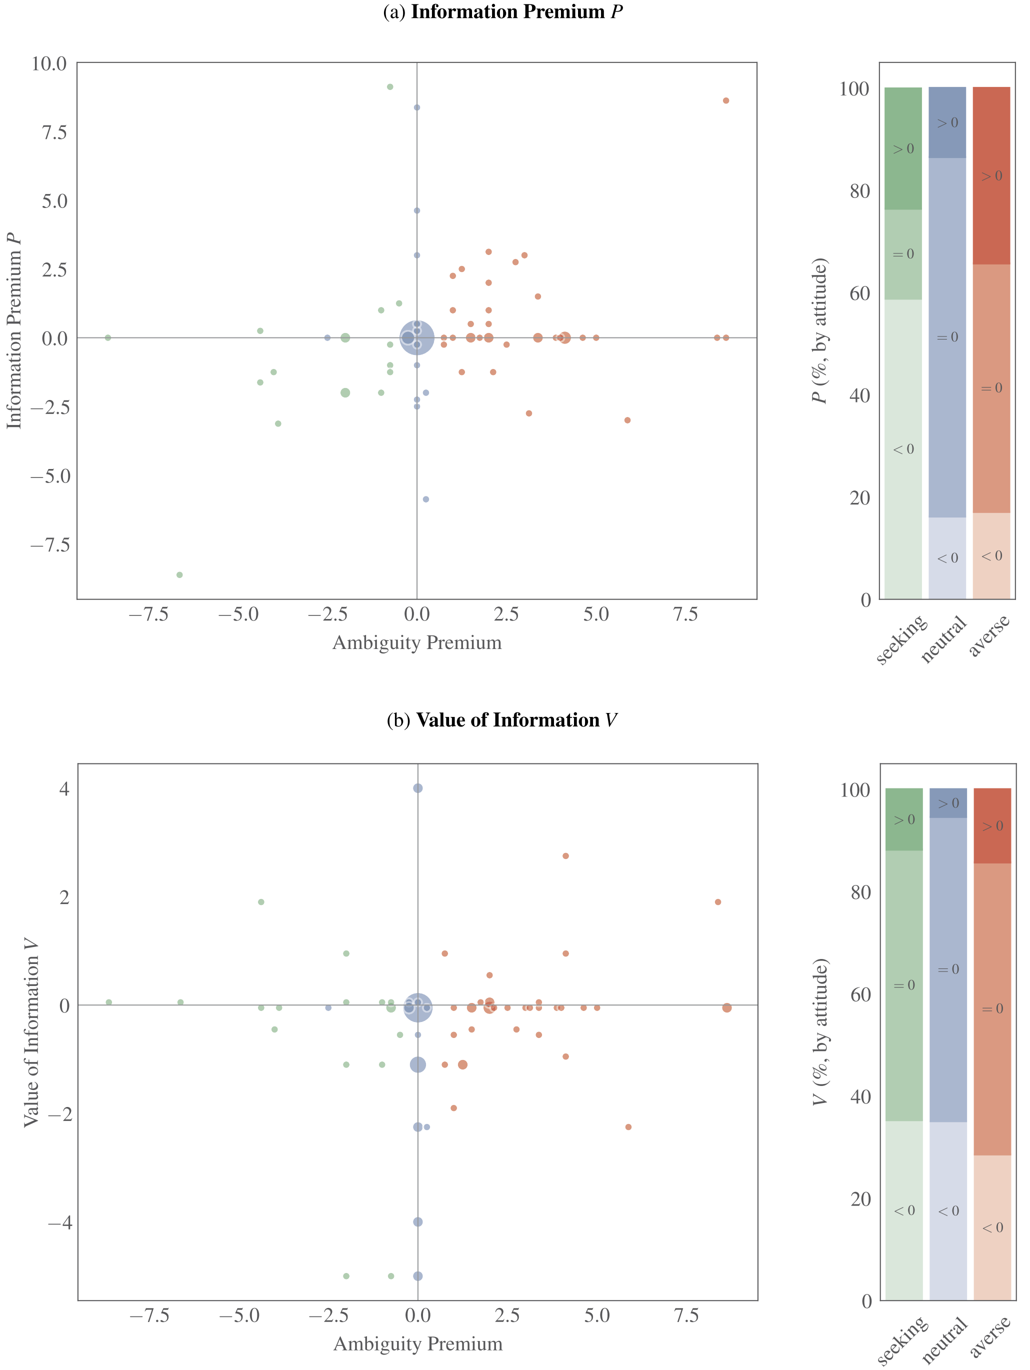

Fig. 2. Results, ambiguous payoff-relevant state, graphically.

3.4. Comparison with theory

Common theoretical models predict that, with risky state, ambiguity-averse individuals should have *negative* Information Premium and Value of Information. Instead, we find that the majority has zero for both. Only a minority (31%) has a negative Information Premium. For ambiguity-seeking subjects, a large majority has a positive Information Premium, as predicted by the models. However, this is not reflected in the Value of Information. While theories predict it should be positive, it is too often zero.

3.5. Concerns

Noise in Q1 As mentioned above, the answer to Q1 is used to compute both the Ambiguity and the Information Premia. Noise in this measure has two effects.

First, it may induce a spurious correlation. We have seen that, for ambiguity-averse subjects, we do not have such a correlation and thus do not have this concern. For ambiguity-seeking, however, we do—and this concerns suggest caution in interpreting it.

Second, it may lead to a misclassification of subjects' ambiguity attitude. Suppose $c_{\varnothing}^{\text{observed}} = c_{\varnothing}^{\text{true}} + \varepsilon$. If the realization of ε is negative, this biases P upwards and increases ambiguity *seeking*. We may be misclassifying some individuals as ambiguity-seeking and overestimating their P. Again, this suggests caution in interpreting positive values of P for them.

If the realization of ε is positive, this leads to underestimation of the Information Premium and overestimation of the Ambiguity Premium. But this is not our concern—compared to the theory, we find values of P that are too *high* for ambiguity-averse individuals. Therefore, while noise in elicitation may induce errors, it cannot lead to our main result that the Information Premium is not negative for ambiguity-averse individuals—it instead pushes in the opposite direction.

In general, even accounting for noise in all measures and in how subjects are classified, the main theoretical prediction remains that Information Premia P should be negative for a sizable fraction of the population—ambiguity-averse subjects, typically the majority. However, this is not what we find.

Other forms of noise Noise in the answers to other questions, if independent and with zero mean, would wash away and not bias our results. In fact, compared to theory, we find values with and without information to be too often identical—pointing to consistency, rather than noise.

Complexity Questions with information are more complex, which may add a confound, especially since reactions to complexity are known to relate to ambiguity attitude (Halevy, 2007; Dean and Ortoleva, 2019). But following this literature, complexity should *lower* the Information Premium for ambiguity-averse individuals. Our finding, instead, is that it is too high and thus does not seem to be caused by the confound.

3.6. Discussion

We study ambiguity of information of one particular form: whether the message is truthful or misleading. While only a special case of ambiguous information, it allows us to test an implication of the dilation property of updating models. Common models predict that ambiguity-averse individuals should lower their value of bets after information. That is, for any message, 'all news is bad news.' We then test it and reject it: we find that the large majority of ambiguity-averse individuals do *not* react to this information.

We now turn to discuss possible explanations, including alternative models.

Choosing if to process information and DC updating Recall that, in our experiments, messages do not seem to be generally ignored. First, during the experiment, subjects were required to acknowledge them (clicking on the corresponding color). Second, many subjects did react to information in the sense that their P is non-zero. In particular, this is true for most ambiguity-seeking subjects with risky payoff-relevant urns, as well as many ambiguity-averse ones with ambiguous payoff-relevant urns. Overall, it is plausible that messages are ignored when harmful—by ambiguity-averse subjects with risky payoff-relevant states—and not ignored when beneficial—by ambiguity-seeking ones. In particular, 63% and 54% of ambiguity-averse subjects, respectively, chose to bet on the color from the message in Q2 and Q5, respectively.[19]

A natural interpretation is that individuals choose strategically when and if to process the information before applying any updating. They may simply ignore it; when they don't, they may apply a rule like FB or ML. Note that ignoring information is never useful in the Expected-Utility framework, wherein information has weakly positive value. But this is no longer the case under ambiguity. It may thus be reasonable for subjects to disregard information when harmful—when 'all news is bad news.' Crucially, this approach is outside the FB or ML models. Furthermore, accounting for it may change the implications in applications.

To our knowledge, the only updating rule with similar implications is Dynamically Consistent updating (DC; Hanany and Klibanoff, 2007, 2009). Under DC, before information, agents make choices contingent on each message to maximize the *ex-ante* overall utility; the updating is such that they then want to implement them after information. Applying it to our experiment with risky state, ambiguity-averse agents do not react to information, while ambiguity-seeking ones do. These are their ex-ante optimal choices because the former wants to reduce exposure to ambiguity, while the latter wants to increase it. Such behavior is in line with our findings.

However, DC violates consequentialism: updated beliefs may be influenced by unrealized parts of the decision problem. This may be considered unappealing.[20]

Proxy updating Alternatively, subjects may be following Proxy updating, introduced in Gul and Pesendorfer (2021), precisely with the goal of avoiding the case of dilation after every message; indeed, the motivation includes examples reminiscent of our experiment. Unfortunately, their updating rule cannot be applied to our case because it is defined in the Choquet Expected Utility framework only for capacities satisfying total monotonicity (which is violated in our risky payoff-state case, see Appendix A). However, our results are aligned with the main idea underlying this approach.

Contraction updating Another possibility is that subjects update beliefs via Contraction Updating (Tang, 2022). Under this rule, an MMEU decision-maker's ambiguity may or may not be

[19] At the same time, 44% and 50% of ambiguity-seeking agents chose to bet on the color from the message in Q2 and Q5, respectively. Note, however, that any choice of color is compatible with leading models—ambiguity-seeking agents with symmetric sets of beliefs would be indifferent.

[20] FB and ML satisfy consequentialism but violate dynamic consistency. The two properties generally conflict under ambiguity (Siniscalchi, 2009). Dominiak et al. (2012) and Bleichrodt et al. (2021) test both and find more support for consequentialism.

resolved, depending on the degree of ambiguity of the realized event. However, ambiguous information may never increase ambiguity about a payoff-relevant state. Thus, in our setting, there would never be dilation, and the possibility of 'all news is bad news' would be ruled out. Hence, this is compatible with our main finding that valuations are unchanged after information. See Tang (2022) for more discussion.

Attitude towards ambiguity in information vs. in payoffs As we discussed above, our interpretation relies on the assumption that the attitude towards ambiguity is the same for states that determine payoffs and information.[21] This is further supported by the use, in the experiment, of the same description for ambiguous payoff-jars and ambiguous information-bags. At the same time, our findings are compatible with the possibility that subjects are payoff-ambiguity averse or seeking, but information-ambiguity neutral.

Implications Our results are incompatible with the dilation property implied by widespread models of updating under ambiguity. Some may view this as unsurprising—especially those who deemed this property unappealing in the first place. Others may note that parsimonious models are necessarily inaccurate, and the issue then becomes the importance of these violations.

Our experiment studies a stark, extreme case that may be unlikely to occur in natural settings. However, it provides the opportunity to directly test a property that has relevant and direct implications in applications, especially in strategic settings. Our results suggest some caution in adopting widespread models of updating under ambiguity, which may be less appealing when information is ambiguous, and towards a consideration of a broader class of models that better account for this case.

Data availability

The data and code are available at http://doi.org/10.5281/zenodo.7596220.

Appendix A. On proxy updating

Gul and Pesendorfer (2021) introduce the Proxy updating rule with the goal of addressing the possibility of 'all news is bad news.' However, this is currently defined only for totally monotone capacities; unfortunately, we cannot express our preferences this way, at least when the state space is risky. To see why, consider the framework introduction in Section 2 for the case in which the payoff-relevant state is risky. Our assumptions for this case are that we have a set of beliefs $\Pi \subseteq \Delta(\Omega \times M)$ such that for each $\pi \in \Pi$:

$$\pi(R) = \pi(B) = 0.5, \quad \pi(r|R) = \pi(b|B), \quad \pi(b|R) = \pi(r|B).$$

Any set of beliefs with these characteristics does not induce a totally monotone capacity. Note that the conditions above imply $\pi(R,r) = \pi(B,b)$ and $\pi(B,r) = \pi(R,b)$. Note also that $\pi(R) = \pi(R,r) + \pi(R,b) = \pi(R,r) + \pi(B,r) = \pi(r) = 0.5$. Similarly we obtain $\pi(b) = 0.5$. Let ρ denote the capacity induced by Π. We know

$$\rho(R) + \rho(B) = \rho(r) + \rho(b) = 1.$$

[21] More precisely, we assume that, if they are ambiguity averse (loving) for states that determine payoffs, they are so also for states that determine the information, although the degrees of this aversion (loving) may vary.

Suppose that ρ is totally monotone. Then its Möbius transform λ must satisfy

$$\lambda(R,r)+\lambda(R,b)+\lambda(B,r)+\lambda(B,b)+\lambda(R)+\lambda(B)=1 \tag{1}$$
$$\lambda(R,r)+\lambda(R,b)+\lambda(B,r)+\lambda(B,b)+\lambda(r)+\lambda(b)=1. \tag{2}$$

However, (1) implies that $\lambda(r)+\lambda(b)=0$. Combining with (2), it implies

$$\lambda(R,r)+\lambda(R,b)+\lambda(B,r)+\lambda(B,b)=1,$$

i.e., there is no ambiguity, and this only happens when $|\Pi|=1$, contradiction.

Appendix B. Examples and proofs

As a preliminary result, we prove a useful lemma that shows that if the sign of the Information Premium is the same for all messages, the Value of Information must have the same sign.

Lemma 1. *Consider an agent whose preferences follow either MMEU or the Smooth Model. Then, if both P_r and P_b are positive (negative, zero, respectively), then V is positive (negative, zero, respectively).*

Proof. First, note that since u is continuous and strictly increasing, c_m is continuous and strictly increasing in the first argument.

Second, fix any $x, y \in \mathbb{R}$ and suppose $c_m(x,y)-c_\varnothing(x,y)=P_m>0$ for each $m \in \{r,b\}$. For each $m \in M$, the map $h_m : V \mapsto u(c_m(x-V,y))-u(c_\varnothing(x,y))$ is continuous, strictly decreasing, positive at 0 and negative at y. Call a function *nice* if it satisfies these properties.

Now note that for any $\pi \in \Delta(M)$, $h_\pi := \pi(r)h_r + \pi(b)h_b$ is also nice. For any $\phi : \mathbb{R} \to \mathbb{R}$, strictly increasing and continuous, define $F^\phi(t) := \phi(t+u(c_\varnothing(x,y))) - \phi(u(c_\varnothing(x,y)))$. Since F_ϕ is continuous, strictly increasing, and sign-preserving (i.e., $F_\phi(0)=0$), then $F_\phi \circ h_\pi$ is also nice. It follows that $h_{\min} := \min_{\pi\in\Pi^M} h_\pi$, $h_{\max} := \max_{\pi\in\Pi^M} h_\pi$, and $h_{\phi,\mu} := \int F_\phi \circ h_\pi \,\mathrm{d}\mu$ are also nice. Since this implies that they are strictly decreasing, positive at 0, and negative at y, it follows that each of these functions must have a unique root which is positive.

Finally, note that V is defined as the root of one of these functions, depending on the model and on the ambiguity attitude. This concludes the proof for the case of $P_m > 0$ for each m.

The case of $P_m = 0$ is trivial and the case of $P_m < 0$ is the equivalent, except that all relevant functions are negative at 0 instead of positive. ||

B.1. Examples of dilation and contraction with ambiguous states

Below are examples in which the set of beliefs can contract, dilate, and remain unchanged with ambiguous information when the payoff-relevant state is ambiguous.

Example 1 *(Contraction).* Fix any $a \in (0, 0.5)$, and let $\Pi = \mathrm{co}(\pi_1, \pi_2)$, where

$$\pi_1^\Omega(R)=a, \qquad \pi_1(r|R)=\pi_1(b|B)=1-a,$$
$$\pi_2^\Omega(R)=1-a, \qquad \pi_2(r|R)=\pi_2(b|B)=a.$$

Here the shape of Π induces a 'negative correlation:' for each $\pi \in \Pi$, the Bayesian posterior on R after message r is 0.5. There is no more ambiguity. Thus, with both FB and ML, $\Pi_r^\Omega = \{0.5\} \subset [a, 1-a] = \Pi^\Omega$. It follows that ambiguity-averse agents have $P_r > 0$ and ambiguity-seeking $P_r < 0$.

Example 2 *(Dilation).* Fix any $\varepsilon \in (0, 0.5)$, let $\Pi = \text{co}(\pi_{00}, \pi_{01}, \pi_{10}, \pi_{11})$, where

$$\begin{aligned}
\pi_{00}^{\Omega}(R) &= \varepsilon, & \pi_{00}(r|R) &= \pi_{00}(b|B) = \varepsilon, \\
\pi_{01}^{\Omega}(R) &= \varepsilon, & \pi_{01}(r|R) &= \pi_{01}(b|B) = 1-\varepsilon, \\
\pi_{10}^{\Omega}(R) &= 1-\varepsilon, & \pi_{10}(r|R) &= \pi_{10}(b|B) = \varepsilon, \\
\pi_{11}^{\Omega}(R) &= 1-\varepsilon, & \pi_{11}(r|R) &= \pi_{11}(b|B) = 1-\varepsilon.
\end{aligned}$$

Intuitively, Π includes all combinations of marginals over Ω and over whether the message is informative or misleading, as if obtained as the 'product' of the two sets. In this case, we have

$$^{\text{FB}}\Pi_m^{\Omega} = {}^{\text{ML}}\Pi_r^{\Omega} = \left[\frac{\varepsilon^2}{\varepsilon^2 + (1-\varepsilon)^2}, \frac{(1-\varepsilon)^2}{\varepsilon^2+(1-\varepsilon)^2}\right] \supsetneq [\varepsilon, 1-\varepsilon] = \Pi^{\Omega}.$$

Thus, under FB and ML, $P_m < 0$ and $V < 0$ if the agent is ambiguity-averse; $P_m > 0$ and $V > 0$ if ambiguity-seeking.

Example 3 *(Unchanged).* Let $\Pi = \text{co}(\pi_{00}, \pi_{01}, \pi_{10}, \pi_{11})$ where π_{ij} are defined as in the previous example, but with $\varepsilon = 0$. Then we have

$$^{\text{FB}}\Pi_r^{\Omega} = {}^{\text{ML}}\Pi_r^{\Omega} = [0, 1] = [0, 1] = \Pi^{\Omega}.$$

Thus, independently of the ambiguity attitude we have $P_m = 0$ and $V = 0$.

B.2. Proofs

Proof of Proposition 1. Notice first that, since u is strictly increasing and $X = \mathbb{R}$, we have that P_r, P_b and V are both well-defined and unique.

Consider first the case of ambiguity-neutrality. Recall that we have assumed $\pi(r|R) = \pi(b|B)$ and $\pi(b|R) = \pi(r|B)$ and that, if $\bar{\pi}$ is the belief for which messages are not informative—$\bar{\pi}(r|R) = \bar{\pi}(b|R) = 0.5$—we have $\bar{\pi} \in \Pi$. When $\Pi = \{\pi\}$, we must then have $\pi(R, r) = \pi(R, b) = \pi(B, b) = \pi(B, r) = 0.25$. In turns, this implies that the decision-maker's belief over R and B will not change after receiving message r or b. Item (3) of the Proposition thus holds.

Consider now a set of beliefs Π with $|\Pi| > 1$. Because $\pi(R) = 0.5$ for all $\pi \in \Pi$, then there must exist $\pi_1, \pi_2 \in \Pi$ such that $\pi_1 \neq \pi_2$ and $\pi_1(R, r) + \pi_1(B, b) \neq \pi_2(R, r) + \pi_2(B, b)$. Denote $\bar{p} := \max_{\pi \in \Pi} \pi(R, r) + \pi(B, b)$ and $\underline{p} := \min_{\pi \in \Pi} \pi(R, r) + \pi(B, b)$. Our assumptions on Π imply $\bar{p} > 0.5 > \underline{p}$.

Assume now the agent is ambiguity-averse and that the updating rule is FB. After message r, the Bayesian update of $\pi \in \Pi$ is $\pi(R|r) = 2\pi(R, r) = \pi(R, r) + \pi(B, b)$. Thus, $\min_{\pi \in \Pi} \pi(R|r) = \underline{p} < 0.5$ and $\min_{\pi \in \Pi} \pi(B|r) = 1 - \bar{p} < 0.5$. No matter what color the agent chooses, she is worse off compared to before the message. Then: $P_m < 0$ for each m; by Lemma 1, $V < 0$. The case of message b is identical. The case of ambiguity-seeking follows identical steps, minimally adapted.

Finally, assume that the agent's updating is ML. Note that under any belief π, $\pi(r) = \pi(b) = 0.5$. Therefore, ML updating is exactly the same as the ML updating, and the result maintains. ‖

Proof of Proposition 2. The case for ambiguity-neutrality is the same as Proposition 1.

If $\Pi > 1$, by Example 1 in the main body, P_m can be positive for ambiguity-averse agents and negative for ambiguity-seeking agents when the updating is FB and ML. By Example 2 in

Section B.1, P_m and V can be negative for ambiguity-averse agents and positive for ambiguity-seeking agents when the updating is FB and ML. By Example 3 in Section B.1, P_m and V can be zero for ambiguity-seeking and ambiguity-averse agents when the updating is FB and ML. We are left to show that under both FB and ML it is possible to have $V > 0$ for ambiguity-averse agents; or $V < 0$ for ambiguity-seeking ones.

We show both with the following example. Consider $\Pi = \mathrm{co}(\pi_1, \pi_2)$ such that

$$\pi_1^{\Omega}(R) = \varepsilon, \qquad \pi_1(r|R) = \pi_1(b|B) = 1-\delta,$$
$$\pi_2^{\Omega}(B) = \varepsilon, \qquad \pi_2(b|R) = \pi_2(r|B) = 1-\delta,$$

where $\varepsilon, \delta < 0.5$. It is easy to check that in this case, ML and FB coincide. When the agent is ambiguity-averse, $P_r = 0.5 - \varepsilon$ and $P_b = \frac{\varepsilon\delta}{\varepsilon\delta + (1-\varepsilon)(1-\delta)} - \varepsilon$. Thus, that $V = 0.5\delta - 0.5\varepsilon$. Therefore, $V > 0$ if $\varepsilon < \delta$. When the agent is ambiguity-seeking, $P_r = \varepsilon - 0.5$ and $P_b = \varepsilon - \frac{\varepsilon\delta}{\varepsilon\delta + (1-\varepsilon)(1-\delta)}$. Thus $V < 0$ if $\varepsilon < \delta$. $\quad\|$

Proof of Proposition 3. Without loss of generality, assume that $u(x+y) = 1, u(x) = 0$. For any $\pi \in \operatorname{supp}\mu$ and any $m \in \{r, b\}$, note that

$$\pi(m) = \pi(m|R)\pi(R) + \pi(m|B)\pi(B) = \frac{1}{2} \tag{3}$$

This implies that $\mu(m|\Pi) = \frac{1}{2}$ and, therefore, $\mu(\Pi|m) = \mu(\Pi)$. It follows that for all $m \in \{r, b\}$

$$\mu(\cdot|m) = \mu. \tag{4}$$

Without information, we have

$$u(c_{\varnothing}) = \max_{\omega \in \{R,B\}} U(x + y\mathbb{1}_{\omega}),$$

where $\forall \omega \in \{R, B\}$

$$\begin{aligned} U(x\mathbb{1}_{\omega}) &= \int_{\Delta(S)} \phi\left[\int_S u \circ (x + y\mathbb{1}_{\omega})\, \mathrm{d}\pi\right] \mathrm{d}\mu \\ &= \int_{\Delta(S)} \phi\left[\pi(\omega)u(x+y) + (1-\pi(\omega))u(x)\right] \mathrm{d}\mu \\ &= \phi\left(\tfrac{1}{2}\right) \end{aligned}$$

Similarly, when message m is observed, the utility is

$$u(c_m) = \max_{\omega \in \{R,B\}} U(x + y\mathbb{1}_{\omega}|m).$$

By (4), for all $\omega \in \{R, B\}$,

$$\begin{aligned} U(x + y\mathbb{1}_{\omega}|m) &= \int_{\Delta(S)} \phi\left[\int_S u \circ (x + y\mathbb{1}_{\omega})\, \mathrm{d}\pi(\cdot|m)\right] \mathrm{d}\mu(\cdot|m) \\ &= \int_{\Delta(S)} \phi\left[\pi(\omega|m)\right] \mathrm{d}\mu. \end{aligned}$$

Now we show that the symmetry of μ implies that

$$\int_{\Delta(S)} \pi(\omega|m)\,\mathrm{d}\mu = \frac{1}{2}. \tag{5}$$

To see why, denote

$$\begin{aligned}
\Delta_+ &= \{\pi \in \Delta(S) | \pi(m|\omega) > \tfrac{1}{2}\},\\
\Delta_0 &= \{\pi \in \Delta(S) | \pi(m|\omega) = \tfrac{1}{2}\},\\
\Delta_- &= \{\pi \in \Delta(S) | \pi(m|\omega) < \tfrac{1}{2}\},
\end{aligned}$$

and note that symmetry implies

$$\mu(\Delta_+) = \mu(\Delta_-) = \tfrac{1}{2} - \tfrac{1}{2}\mu(\Delta_0), \quad \mu \circ G_\omega^{-1} = \mu \circ (1 - G_\omega)^{-1},$$

where $G_\omega : \pi \mapsto \pi(\cdot|\omega)$. Therefore, (5) holds because

$$\begin{aligned}
\int_{\Delta(S)} \pi(\omega|m)\,\mathrm{d}\mu &= \int_{\Delta(S)} \frac{\frac{1}{2}\pi(m|\omega)}{\frac{1}{2}\pi(m|R) + \frac{1}{2}\pi(m|B)}\,\mathrm{d}\mu\\
&= \int_{\Delta(S)} \pi(m|\omega)\,\mathrm{d}\mu\\
&= \int_{\Delta_+} \pi(m|\omega)\,\mathrm{d}\mu + \int_{\Delta_-} \pi(m|\omega)\,\mathrm{d}\mu + \int_{\Delta_0} \pi(m|\omega)\,\mathrm{d}\mu\\
&= \int_{\Delta_+} \pi(m|\omega)\,\mathrm{d}\mu + \int_{\Delta_+} (1 - \pi(m|\omega))\,\mathrm{d}\mu + \frac{1}{2}\mu(\Delta_0)\\
&= \mu(\Delta_+) + \frac{1}{2}\mu(\Delta_0) = \frac{1}{2}.
\end{aligned}$$

In the case of ambiguity-averse agent (ϕ strictly concave), by Jensen's inequality,

$$U(x + y\mathbb{1}_\omega|m) = \int_{\Delta(S)} \phi\left[\pi(\omega|m)\right]\mathrm{d}\mu < \phi\left[\int_{\Delta(S)} \pi(\omega|m)\,\mathrm{d}\mu\right] = \phi\left[\tfrac{1}{2}\right] = u(c_\varnothing).$$

To sum up, we have $\forall m \in \{r, b\}, \omega \in \{R, B\}, U(x + y\mathbb{1}_\omega|m) < u(c_\varnothing)$, which implies $c_m < c_\varnothing$ and, therefore, $P_m < 0$, $V < 0$.

The case of affine ϕ is straightforward and the case of strictly convex ϕ is analogous to the one above. ‖

Appendix C. Additional experimental analysis

C.1. Additional tables

Table 3 summarizes ambiguity attitude in the sample; Table 4 summarizes ambiguity premia in the sample; Table 5 is similar to Table 2, but with statistics grouped by order.

Table 3
Ambiguity attitude in the sample.

	averse	neutral	seeking	Total
A	23	22	7	52
B	12	15	10	37
Total	35	37	17	89

Table 4
Ambiguity premia in the sample.

amb. att.	averse			neutral			seeking		
order	A	B	All	A	B	All	A	B	All
median	2.8	2	2.5	0	0	0	−4	−1.5	−2
mean	3.3	2.8	3.1	−0.05	−0.02	−0.03	−3.9	−1.8	−2.7
# of obs.	23	12	35	22	15	37	7	10	17

C.2. Analysis with all subjects

The analysis in the main body of the paper does not include the behavior of two subjects who reported extreme answers in a number of questions—-in particular, they chose a bet on the urn with a payment of $20 against any amount of money, including $20 for sure. Below we replicate Table 3 and Table 2, but include these two subjects. As can be seen from Table 6, both subjects faced Order B, one of them was ambiguity-seeking, another – ambiguity-neutral. Thus, there is very little difference between Tables 2 and 7.

Table 5
Results (by order).

	Risky payoff state								*Ambiguous payoff state*							
amb. attitude	All		averse		neutral		seeking		All		averse		neutral		seeking	
order	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B
								Information Premium P								
theory prediction			< 0	< 0	$= 0$	$= 0$	> 0	> 0			-	-	$= 0$	$= 0$	-	-
median	0	0	0	0	0	0	3.5	0	0	0	0	0.25	0	0	0	−1.1
mean	0.57	0.29	−0.76	−0.17	0.5	0.75	5.2	0.16	−0.17	0.54	0.25	1.2	−0.23	0.53	−1.3	−0.21
≈ 0	62%	70%	61%	83%	77%	80%	14%	40%	67%**	46%	61%	42%	86%*	60%	29%	30%
$= 0$	50%	65%	52%	67%	59%	80%	14%	40%	60%*	40%	52%	42%	82%*	53%	14%	20%
> 0	31%	24%	9%	17%	36%	20%	86%*	40%	19%	30%	26%	50%	4%*	27%	43%	10%
< 0	19%	11%	39%	17%	4%	0%	0%	20%	21%	30%	22%	8%	14%	20%	43%	70%
								Value of Information V								
theory prediction			< 0	< 0	$= 0$	$= 0$	> 0	> 0			-	-	$= 0$	$= 0$	-	-
median	−0.05	−0.05	−0.05	−0.05	−0.05	−0.55	0.05	0	−0.05	−0.05	−0.05	−0.05	−0.05	−0.55	0.05	−0.05
mean	−0.34	−0.52	−0.44	0.17	−0.48	−1.1	0.46	−0.48	−0.28	−0.63	−0.19	0.04	−0.46	−0.9	−0.01	−1
$= 0$	58%	49%	56%	42%	59%	40%	57%	70%	60%	54%	52%	67%	73%**	40%	43%	60%
> 0	8%	14%	4%*	25%	4%	7%	29%	10%	12%	8%	17%	8%	4%	7%	14%	10%
< 0	35%	38%	39%	33%	36%	53%	14%	20%	29%	38%	30%	25%	23%*	53%	43%	30%
# of obs.	52	37	23	12	22	15	7	10	52	37	23	12	22	15	7	10

† Due to rounding, the percentages for > 0, < 0, $= 0$ need not always add up to 100%.
Stars in A-columns indicate (two-sided) statistically significant difference from values in B-columns: $^{*}p < 0.1$, $^{**}p < 0.05$.

Table 6
Ambiguity attitude in the sample (all observations).

	averse	neutral	seeking	Total
A	23	22	7	52
B	12	16	11	39
Total	35	38	18	91

Table 7
Results (all observations).

	Risky payoff state				*Ambiguous payoff state*			
ambiguity attitude	All	averse	neutral	seeking	All	averse	neutral	seeking
				Information Premium P				
theory prediction		< 0	$= 0$	> 0		-	0	-
median	0	0	0	0.88	0	0	0	-0.62
mean	0.42	-0.56	0.59	2	0.11	0.57	0.04	-0.64
≈ 0	65%	69%	79%	28%	59%	54%	76%	33%
$= 0$	56%	57%	68%	28%	52%	49%	68%	22%
> 0	28%	11%	29%	56%	23%	34%	13%	22%
< 0	16%	31%	3%	17%	25%	17%	18%	56%
				Value of Information V				
theory prediction		< 0	$= 0$	> 0		-	0	-
median	-0.05	-0.05	-0.05	0.05	-0.05	-0.05	-0.05	-0.05
mean	-0.41	-0.23	-0.71	-0.09	-0.42	-0.11	-0.63	-0.57
$= 0$	55%	51%	53%	67%	58%	57%	60%	56%
> 0	10%	11%	5%	17%	10%	14%	5%	11%
< 0	35%	37%	42%	17%	32%	29%	34%	33%
# of obs.	91	35	38	18	91	35	38	18

† Due to rounding, the percentages for > 0, < 0, $= 0$ need not always add up to 100%.

Appendix D. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jet.2023.105610.

References

Azrieli, Yaron, Chambers, Christopher P., Healy, Paul J., 2018. Incentives in experiments: a theoretical analysis. J. Polit. Econ. 126 (4), 1472–1503.
Bade, Sophie, 2015. Randomization devices and the elicitation of ambiguity-averse preferences. J. Econ. Theory 159, 221–235.
Baillon, Aurélien, Bleichrodt, Han, Keskin, Umut, L'Haridon, Olivier, Li, Chen, 2017. The effect of learning on ambiguity attitudes. Manag. Sci. 64 (5), 2181–2198.
Beattie, Jane, Loomes, Graham, 1997. The impact of incentives upon risky choice experiments. J. Risk Uncertain. 14 (2), 155–168.
Beauchêne, Dorian, Li, Jian, Li, Ming, 2019. Ambiguous persuasion. J. Econ. Theory 179, 312–365.
Bleichrodt, Han, Eichberger, Jürgen, Grant, Simon, Kelsey, David, Li, Chen, 2021. Testing dynamic consistency and consequentialism under ambiguity. Eur. Econ. Rev. 134, 103687.

Bose, Subir, Renou, Ludovic, 2014. Mechanism design with ambiguous communication devices. Econometrica 82 (5), 1853–1872.
Chapman, Jonathan, Dean, Mark, Ortoleva, Pietro, Snowberg, Erik, Camerer, Colin, forthcoming. Econographics. J. Polit. Econ. Microecon.
Chen, Daniel L., Schonger, Martin, Wickens, Chris, 2016. OTree—an open-source platform for laboratory online, and field experiments. J. Behav. Exp. Finance 9, 88–97.
Cheng, Xiaoyu. Ambiguous Persuasion under Dynamic Consistency. Florida State University. 2021. Mimeo.
Cohen, M., Gilboa, I., Jaffray, J.Y., Schmeidler, D., 2000. An experimental study of updating ambiguous beliefs. Risk Decis. Policy 5 (2), 123–133.
Cubitt, Robin, Starmer, Chris, Sugden, Robert, 1998. On the validity of the random lottery incentive system. Exp. Econ. 1 (2), 115–131.
De Filippis, Roberta, Guarino, Antonio, Jehiel, Philippe, Kitagawa, Toru, 2022. Non-Bayesian updating in a social learning experiment. J. Econ. Theory 199, 105188.
Dean, Mark, Ortoleva, Pietro, 2019. The empirical relationship between nonstandard economic behaviors. Proc. Natl. Acad. Sci. 116 (33), 16262–16267.
Dempster, A.P., 1967. Upper and lower probabilities induced by a multivalued mapping. Ann. Math. Stat. 38 (2), 325–339.
Dominiak, Adam, Duersch, Peter, Lefort, Jean-Philippe, 2012. A dynamic Ellsberg Urn experiment. Games Econ. Behav. 75 (2), 625–638.
Epstein, Larry, Halevy, Yoram. Hard-to-Interpret Signals. University of Toronto. 2022. Mimeo.
Epstein, Larry G., Schneider, M., 2007. Learning under ambiguity. Rev. Econ. Stud. 74 (4), 1275–1303.
Epstein, Larry G., Schneider, Martin, 2008. Ambiguity, information quality, and asset pricing. J. Finance 63 (1), 197–228.
Ert, Eyal, Trautmann, Stefan T., 2014. Sampling experience reverses preferences for ambiguity. J. Risk Uncertain. 49 (1), 31–42.
Esponda, Ignacio, Vespa, Emanuel. Contingent preferences and the sure-thing principle: revisiting classic anomalies in the laboratory. UCSB. 2021. Mimeo.
Freeman, David, Halevy, Yoram, Kneeland, Terri, 2019. Eliciting risk preferences using choice lists. Quant. Econ. 10, 217–237.
Ghirardato, Paolo, Maccheroni, Fabio, Marinacci, Massimo, 2008. Revealed ambiguity and its consequences: updating. In: Abdellaoui, J., Hey, Mohammedand (Eds.), Advances in Decision Making Under Risk and Uncertainty. Selected Papers from the FUR 2006 Conference. Springer-Verlag, Berlin.
Gilboa, Itzhak, Schmeidler, David, 1989. Maxmin expected utility with non-unique prior. J. Math. Econ. 18, 141–153.
Gilboa, Itzhak, Schmeidler, David, 1993. Updating ambiguous beliefs. J. Econ. Theory 59, 34–49.
Gilboa, Itzhak, Marinacci, Massimo, 2013. Ambiguity and the Bayesian paradigm. In: Acemoglu, Daron, Arellano, Manuel, Dekel, Eddie (Eds.), Advances in Economics and Econometrics: Theory and Applications, Tenth World Congress of the Econometric Society. Cambridge University Press.
Good, I.J., 1967. On the principle of total evidence. Br. J. Philos. Sci. 17 (4), 319–321.
Good, I.J., 1974. A little learning can be dangerous. Br. J. Philos. Sci. 25 (4), 340–342.
Gul, Faruk, Pesendorfer, Wolfgang, 2021. Evaluating ambiguous random variables from Choquet to maxmin expected utility. J. Econ. Theory 192, 105129.
Halevy, Y., 2007. Ellsberg revisited: an experimental study. Econometrica 75 (2), 503–536.
Hanany, Eran, Klibanoff, Peter, 2007. Updating preferences with multiple priors. Theor. Econ. 2 (3), 261–298.
Hanany, Eran, Klibanoff, Peter, 2009. Updating ambiguity averse preferences. B. E. J. Theor. Econ. 9 (1), 37.
Hey, John, Lee, Jinkwon, 2005. Do subjects separate (or are they sophisticated)? Exp. Econ. 8 (3), 233–265.
Jaffray, J., 1992. Bayesian updating and belief functions. IEEE Trans. Syst. Man Cybern. 22 (5), 1144–1152.
Karni, E., Safra, Z., 1987. 'Preference reversal' and the observability of preferences by experimental methods. Econometrica 55 (3), 675–685.
Kellner, Christian, Le Quement, Mark T., Riener, Gerhard, 2022. Reacting to ambiguous messages: an experimental analysis. Games Econ. Behav. 136, 360–378.
Klibanoff, Peter, Marinacci, Massimo, Mukerji, Sujoy, 2005. A smooth model of decision making under ambiguity. Econometrica 73 (6), 1849–1892.
Klibanoff, Peter, Marinacci, Massimo, Mukerji, Sujoy, 2009. Recursive smooth ambiguity preferences. J. Econ. Theory 144 (3), 930–976.
Kops, Christopher, Pasichnichenko, Illia. Testing Negative Value of Information and Ambiguity Aversion. University of Heidelberg. 2022. Mimeo.
Kurata, Hiroshi, Izawa, Hiroshi, Okamura, Makoto, 2009. Non-expected utility maximizers behave as if expected utility maximizers: an experimental test. J. Socio-Econ. 38 (4), 622–629.

Liang, Yucheng. Learning from Unknown Information Sources. Stanford University. 2022. Mimeo.
Moreno, Othon M., Rosokha, Yaroslav, 2016. Learning under compound risk vs. learning under ambiguity–an experiment. J. Risk Uncertain. 53 (2–3), 137–162.
Pahlke, Marieke. Dynamic Consistency and Ambiguous Communication. Paris School of Economics. 2022. Mimeo.
Pires, Cesaltiona Pacheco, 2002. A rule for updating ambigous beliefs. Theory Decis. 53, 137–152.
Seidenfeld, Teddy, 1981. Levi on the dogma of randomization in experiments. In: Henry E. Kyburg, Jr. & Isaac Levi. Springer, pp. 263–291.
Seidenfeld, Teddy, Wasserman, Larry, 1993. Dilation for sets of probabilities. Ann. Stat. 21 (3), 1139–1154.
Shafer, Glenn, 1976. A Mathematical Theory of Evidence. Princeton University Press, Princeton, NJ.
Siniscalchi, Marciano, 2009. Two out of three ain't bad: a comment on "The ambiguity aversion literature: a critical assessment". Econ. Philos. 25 (03), 335–356.
Tang, Rui. A Theory of Contraction Updating. Hong Kong University of Science and Technology. mimeo.
Vinogradov, Dmitri, Makhlouf, Yousef, 2020. Signaling probabilities in ambiguity: who reacts to vague news? Theory Decis. 90, 371–404.
Walley, Peter, 1991. Statistical Reasoning with Imprecise Probabilities. Chapman and Hall, London.
Wasserman, Larry A., Kadane, Joseph B., 1990. Bayes' theorem for Choquet capacities. Ann. Stat. 18 (3), 1328–1339.