Decision

Information Gaps: A Theory of Preferences Regarding the Presence and Absence of Information

Russell Golman and George Loewenstein

Online First Publication, October 10, 2016. http://dx.doi.org/10.1037/dec0000068

Decision

2325-9965/16/$12.00 http://dx.doi.org/10.1037/dec0000068

Information Gaps: A Theory of Preferences Regarding the Presence and Absence of Information

Russell Golman and George Loewenstein
Carnegie Mellon University

We propose a theory of preferences for acquiring or avoiding information and for exposure to uncertainty (i.e., risk or ambiguity) which is based on thoughts and feelings about information as well as information gaps, that is, specific questions that the decision maker recognizes and is aware of. In our theoretical framework utility depends not just on material payoffs but also on beliefs and the attention devoted to them. We specify assumptions regarding the determinants of attention to information gaps, characterize a specific utility function that describes feelings about information gaps, and show with examples that our theory can make sense both of the acquisition of noninstrumental information and of the avoidance of possibly useful information, as well as source-specific risk and ambiguity aversion and seeking.

Keywords: belief-based utility, information gap, uncertainty

Thomas Schelling's characterization of *The Mind as a Consuming Organ* (Schelling, 1987) highlights the fact that most if not all consumption is "in the mind." Schelling's observation presents a challenge for the revealed preference philosophy so prevalent in economics. We cannot directly observe the objects of preference consumed in the mind.[1] What is the mind consuming (or preferring not to consume) when we observe people succumbing to clickbait on the Internet, or skipping a visit to the doctor despite unrelieved symptoms of illness, or gambling on their favorite sports teams after purchasing a low-deductible insurance policy? Distinct behavioral economic models, based on psychology rather than revealed preference, can account for some of these patterns of behavior, but none provides an integrated explanation of the diversity of information-related behaviors. Here, we develop a theory that generates predictions about when people will obtain or avoid information and when they will exhibit risk and ambiguity seeking or aversion, based on psychologically grounded assumptions about how people think and feel about the presence and absence of information.

At the core of our theory is the concept of an *information gap* (Loewenstein, 1994). There are many things that one does not know and does not think about, but when a person is aware of a specific unknown, it often attracts attention and evokes emotions. We introduce a reduced form question-answer framework for representing knowledge and awareness. In this framework, an information gap opens when a person becomes aware of a question and is uncertain about the correct answer. People form beliefs about information gaps, making judgments about the chance that each answer is correct. We make a series of assumptions describing how much attention each of these beliefs attracts, and we propose a specific utility function defined over these beliefs and the attention paid to them.

Russell Golman and George Loewenstein, Department of Social and Decision Sciences, Carnegie Mellon University.

We thank Botond Köszegi, David Laibson, Erik Eyster, Juan Sebastian Lleras, Alex Imas, Leeann Fu, and Donald Hantula for helpful comments on an earlier draft entitled *Curiosity, Information Gaps, and the Utility of Knowledge*.

Correspondence concerning this article should be addressed to Russell Golman, Department of Social and Decision Sciences, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213. E-mail: rgolman@andrew.cmu.edu

[1] Just as psychology has moved beyond behaviorism and embraced cognition, economics also has much to gain by acknowledging the inner workings of the mind (Bernheim & Rangel, 2009; Chater, 2015; Kimball, 2015). Caplin and Leahy (2004), for example, offer a clear demonstration that the revealed preference framework is insufficient to analyze benevolent information disclosure policies.

From these primitives, we derive predictions about how people will respond to information gaps, in terms of whether they will seek or avoid information, risk, and ambiguity.

Relation to Existing Literature

Background

The study of decision making under uncertainty grew out of expected utility theory (Anscombe & Aumann, 1963; Savage, 1954; von Neumann & Morgenstern, 1944). Traditionally, choices were thought of in consequentialist terms, with outcomes corresponding to objective reality (Hammond, 1988). According to expected utility theory, risk preferences could be described by utility function curvature, and people would not be sensitive to ambiguity.

The first economic analysis of information preference was pioneered by Stigler (1961). Drawing attention to a phenomenon that had hitherto not been addressed by economists, Stigler assumed that information is a means to an end; it is valued because, and only to the extent that, it enables people to make better decisions. Preferences about information could then be derived from expected utility theory, assuming Bayesian updating of beliefs, with the implication that information will be sought to the degree that it raises expected utility (cf. Hirshleifer & Riley, 1979). According to Stigler's theory, information could never have negative value; people would never deliberately avoid receiving information (in private).

Of course, people do not generally conform to expected utility theory or, in part for that reason, value information in accordance with Stigler's theory. Drawing on a set of psychologically grounded assumptions, Kahneman and Tversky (1979; Tversky & Kahneman, 1992) developed prospect theory to account for violations of expected utility, such as the simultaneous purchase of insurance and lottery tickets, and economists developed more general theories of decision making under uncertainty to represent a range of possible preferences about risk and ambiguity, based (unlike prospect theory) on a revealed preference approach (e.g., Abdellaoui et al., 2011; Gul, 1991; Klibanoff et al., 2005). Nonexpected-utility theories led economists to revisit the value of information.

Given exposure to some uncertainty, acquiring information about that uncertainty can be seen as resolving a lottery over the remaining uncertainty. The value of this information would be the difference between the utility of the compound lottery and the utility of the original prospect, and for nonexpected-utility theories, this value could be negative (Andries & Haddad, 2015; Dillenberger, 2010; Grant et al., 1998; Kreps & Porteus, 1978; Wakker, 1988). Such theories can thus describe certain instances of information avoidance.

Theories of belief-based utility represent another step in the progression of economists' analysis of information preferences that began with Stigler. These theories recognize that people derive utility not (only) from objective reality but from their beliefs about that reality, for example, their anticipatory feelings (Loewenstein, 1987).[2] From this perspective, acquiring information can be seen as resolving a lottery about what the person may come to believe. Risk aversion (lovingness) over beliefs implies that people will want to avoid (obtain) information (Caplin & Leahy, 2001; Köszegi, 2003; Schweizer & Szech, 2014). Risk aversion over beliefs (and hence information avoidance) could develop when people hold favorable beliefs and do not want to lose them (e.g., Benabou & Tirole, 2002; Köszegi, 2006) or when people are generally loss averse (e.g., Köszegi, 2010).

Existing belief-based utility theories clearly do make predictions about when people will obtain information and when they will avoid it, and prospect theory clearly does make predictions about when people will seek or steer clear of risk and uncertainty, but some stylized facts in both domains still call out for explanation.

First, (in contrast to the predictions of Benabou and Tirole (2002) and Köszegi (2006)), information avoidance is more common while holding unfavorable beliefs than while holding favorable beliefs (Dwyer et al., 2015; Eil & Rao, 2011; Fantino & Silberberg, 2010; Ganguly & Tasoff, 2016; Lieberman et al., 1997; Karlsson et al., 2009). Existing belief-based utility models, which assume that risk preferences over beliefs are independent of those be-

[2] See also Abelson, 1986; Geanakoplos et al., 1989; Asch et al., 1990; Yariv, 2001.

liefs, cannot explain such a pattern because with Bayesian updating expost beliefs cannot be expected to be better or worse than exante beliefs (Eliaz & Spiegler, 2006).

Second, information acquisition often occurs in situations in which a person does not care what he finds out, for example, answers to trivia questions (Berlyne, 1954, 1960; Kang et al., 2009; Marvin & Shohamy, 2016). Existing models of belief-based utility cannot account for such pure curiosity because when the possible outcomes all have the same utility, the sure-thing principle implies that a mixture between these outcomes also has the same utility. In fact, people sometimes seek out information which they know will make them miserable (Hsee & Ruan, 2016; Kruger & Evans, 2009), which poses an even more fundamental challenge to existing belief-based utility models.

Third, information acquisition or avoidance is highly dependent on situational determinants, such as awareness of related uncertainties, the presence of clues about the information content, or opportunities for distraction (Falk & Zimmermann, 2016; Litman et al., 2005; Loewenstein, 1994; Menon & Soman, 2002; van Dijk & Zeelenberg, 2007). Few, if any, of these patterns are predicted by existing models of information preference.

Stylized facts in the domain of preference under risk and ambiguity also call out for explanation. While people often exhibit risk and ambiguity aversion, they also tend to gamble on uncertainties that they feel they have expertise about (Heath & Tversky, 1991). Additionally, the degree of risk or ambiguity aversion (or seeking) people exhibit depends on contextual factors, such as the presence of other risky or ambiguous options for comparison or for mental accounting (Fox & Tversky, 1995; Gneezy & Potters, 1997). Although existing models (e.g., Abdellaoui et al., 2011) can accommodate these phenomena, they do not specifically predict them.

Our Approach

Here we propose a unified theory that can account for these stylized facts. We follow Caplin and Leahy (2001) and Köszegi (2010) in applying expected utility theory to psychological states rather than to physical prizes, but we expand the domain of psychological states that people can have feelings about. We propose specific assumptions about attention and a specific utility function that takes as inputs beliefs and the attention devoted to them (as well as material payoffs). We incorporate Loewenstein's (1994) insight that information gaps stimulate curiosity, as well as Tasoff and Madarász's (2009) insight that obtaining information stimulates attention and thus complements anticipatory feelings. To the extent that a person is thinking about (i.e., attending to) an information gap, feelings about this information gap enter into utility.[3] We show that our model allows for acquisition of noninstrumental information as well as avoidance of possibly useful information, and that it provides a novel account of ambiguity aversion in the famous Ellsberg paradox along with ambiguity seeking in gambling for pleasure. A companion article (Golman & Loewenstein, 2016) explores the full set of implications of our proposed utility model for information acquisition or avoidance. Another companion article (Golman, Loewenstein, & Gurney, 2016) uses the model developed here to derive and test predictions about risk and ambiguity aversion and seeking. We summarize these analyses in the present article to give the reader a sense of the wide range of stylized facts that our model can reconcile.[4]

In the next section we introduce a framework for representing questions and answers, beliefs about them, and the attention to them. We then provide psychological motivation for a specific utility function that incorporates beliefs and attention, and formally characterize this utility function with seven properties. We proceed to describe how decision making operates, and then present our assumptions about the determinants of attention. In the section on, Information Acquisition and Avoidance, we show how our theory can be applied to make sense of information acquisition due to curiosity as well as information avoidance due to anxiety. In the

[3] Curiosity correlates with brain activity in regions thought to relate to anticipated reward (Kang et al., 2009), suggesting that information is a reward in and of itself. Similarly, making decisions while aware of missing relevant information correlates with brain activity in regions thought to relate to fear (Hsu et al., 2005).

[4] We do not claim to reconcile all behavioral patterns in the domains of information preference and risk and ambiguity preference. No theory is perfect.

section on Risk and Ambiguity Preference, we apply our theory to preferences about uncertain gambles, showing that it can account for both the Ellsberg paradox and gambling for pleasure. The examples in these two sections demonstrate the usefulness of our theory. We conclude with a brief discussion of how our assumptions about how people think and feel about information gaps allow us to integrate our accounts of informational preferences and of risk and ambiguity preferences.

Theoretical Framework

Traditional economic theory assumes that utility is a function of consumption bundles or material outcomes, or (perhaps subjective) distributions thereof. Our basic premise is that utility depends not only on such material outcomes but also on beliefs and the attention paid to them. Figure 1 depicts this perspective and illustrates how we construct our framework. At its core are questions and answers, a structure we introduce to represent the information a person has and the information that he is aware he is missing. Questions delineate the issues that a person is aware of. Each has one or more possible answers. The person may or may not know which answer is actually correct. Moving out from the core in this diagram, a person forms beliefs about the answers to the questions he is aware of. He also pays some degree of attention to each of these beliefs/questions. His *cognitive state*, which we shall define as these beliefs and the attention paid to them, is then the object of his preferences.

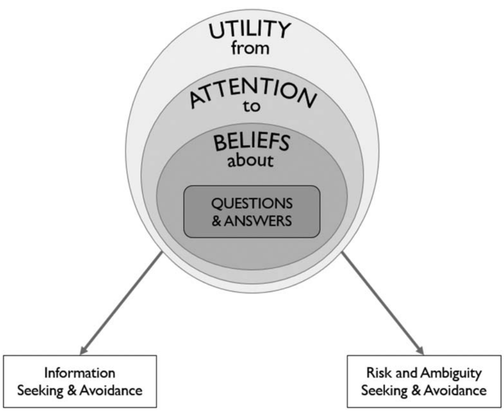

Figure 1. Schematic illustration of the development of our theory.

Cognitive States

While there surely is an infinite set of possible states of the world, we assume that a person can only conceive of a finite number of questions at any one time. We represent awareness with a set of "*activate*"' questions $\mathcal{Q} = \{Q_1, \ldots, Q_m\}$ and a remaining set of "*latent*" questions. Activated questions are those that the individual is aware of (i.e., pays at least some attention to). Latent questions are those that the individual could become, but is not currently, aware of. A vector of attention weights $\mathbf{w} = (w_1, \ldots, w_m) \in \mathbb{R}^m_+$ indicates how much attention each activated question gets.[5]

A question Q_i has a countable set of possible (mutually exclusive) answers $\mathcal{A}_i = \{A_i^1, A_i^2, \ldots\}$.[6] A person has a subjective belief about the probability that each answer is correct. (The subjective probabilities across different questions may well be mutually dependent.) This framework allows us to capture information gaps, which are represented as activated questions lacking known correct answers, as depicted in Table 1.

Anticipated material outcomes, or prizes, can also be incorporated into this framework. We let X denote a countable set of prizes—that is, material outcomes. The subjective probability over these prizes is in general mutually dependent with the subjective probability over answers to activated questions; that is, the receipt of new information often leads to revised beliefs about the likelihood of answers to many different questions as well as about the likelihood of different material outcomes. Denote the space of answer sets together with prizes as $\alpha = \mathcal{A}_1 \times \mathcal{A}_2 \times \cdots \times \mathcal{A}_m \times X$. Then, given a state of awareness defined by the set of activated questions $\mathcal{Q}$,[7] we represent a person's cognitive state C with a probability measure π defined

[5] We can think of the (presumably infinite) set of latent questions as having attention weights of zero.

[6] We use the term countable here to mean *at most countable*. The restriction of a countable set of answers to a countable set of possible questions does still allow an uncountable set of possible states of the world, but as awareness is finite, the precise state of the world would be unknowable.

[7] In most cases, we will assume that activation of questions is determined exogenously—that is, by the environment. We do not model growing awareness (see Karni & Viero, 2013).

Table 1
The Question–Answer Knowledge Structure

Question	Answer	Belief	
Latent	—	Unawareness	
Activated	Unknown	Uncertainty	$\updownarrow$ information gap
	Known	Certainty	

over α (i.e., over possible answers to activated questions as well as eventual prizes) and a vector of attention weights **w**. We denote the set of all possible cognitive states as $\mathcal{C} = \Delta(\alpha) \times \mathbb{R}^m_+$ (with the notation $\Delta(\alpha)$ referring to the space of probability distributions over α with finite entropy).[8] Each marginal distribution π_i specifies the (subjective) probability of possible answers to question Q_i, and similarly π_X specifies the (subjective) probability over prizes.[9]

The formal representation of a cognitive state is depicted in Table 2. Consider, for example, a college professor deciding whether or not to look at her teaching ratings. The set of activated questions (and possible answers) might include: "How many of my students liked my teaching?" (0, 1, 2, . . .); "Did they applaud on the last day of class?" (yes/no); "How good a teacher am I?" (great, good, so-so, bad, awful); "Will I get tenure?" (yes/no). Prior belief about the first question might be quite uncertain. The answer to the second question, on the other hand, might already be known with certainty. There may or may not be much uncertainty about the third and fourth questions. All of these beliefs (to the extent they are uncertain) are jointly dependent. Each of these beliefs also attracts some degree of attention. The material outcome might be next year's salary, which would also be highly, but not perfectly, correlated with whether or not she gets tenure.

Preferences Over (Distributions of) Cognitive States

The conventional theory of choice under risk assumes that a lottery over outcomes is evaluated according to its expected utility. We make the analogous assumptions leading to an expected utility representation for lotteries over cognitive states.

We assume that *there is a complete and transitive preference relation* $\succeq$ *on* $\Delta(\mathcal{C})$ *that is continuous*[10] *and that satisfies independence*, so there exists a continuous expected utility representation u of $\succeq$ (von Neumann & Morgenstern, 1944).

The assumption of *independence across cognitive states* here means that when information could put a person into one of many possible cognitive states, preference is consistent with valuing each possible cognitive state independently of any other cognitive states the person might have found herself in.[11]

The Utility Function

To generate testable predictions, we need to propose a utility function. In this section we specify a utility function that describes feelings about information gaps.

Psychological Insights

We first introduce a few psychological insights about the effects of attention and feelings about uncertainty to motivate our specific utility function.[12] Neuroeconomic research indicates that attention shapes preference and valuation (cf. Fehr & Rangel, 2011). Attention weights in our model specify how much a person is thinking about particular beliefs and, in turn, how much those beliefs directly impact utility. We

[8] The restriction to distributions with finite entropy serves a technical purpose, but it should not trouble us—intuitively, it means that a person cannot be aware of an infinite amount of information, which is also the basis for our assumption that the set of activated questions is finite.

[9] For any $\tilde{\mathcal{A}} \subseteq \mathcal{A}_i$, we have $\pi_i(\tilde{\mathcal{A}}) = \pi(\mathcal{A}_1 \times \cdots \times \mathcal{A}_{i-1} \times \tilde{\mathcal{A}} \times \mathcal{A}_{i+1} \times \cdots \times \mathcal{A}_m \times X)$.

[10] The induced topology on $\mathcal{C}$ (derived from the order topology on $\Delta(\mathcal{C})$) should be a refinement of the order topology on $\mathcal{C}$ (see Nielsen, 1984).

[11] This might seem to imply that the utility of a state of uncertain knowledge is equal to the expected utility of each of the possible beliefs—for example, that being uncertain of whether the object of my desire reciprocates my affections provides the same utility as the sum of probabilities times the utilities associated with the possible outcome belief states. It need not, because (as we discuss in detail below) obtaining the information, and indeed the specific information one obtains, is likely to affect one's attention weights. Such a change in attention can encourage or discourage a decision maker from resolving uncertainty, depending on whether the news that will be revealed is expected to be good or bad.

[12] In a sense we are trying to use a notion of hedonic utility to serve as a decision utility. We do not propose that people actually compute utilities and optimize as part of the decision process, but only that they use some heuristic process that is well adapted to often maximize hedonic utility.

Table 2
Representation of a Cognitive State

Activated questions	Possible answers	Subjective probabilities[a]	Attention weights
Q_1	$\mathcal{A}_1 = \{A_1^1, A_1^2, \cdots\}$	$[\pi_1(A_1^1), \pi_1(A_1^2), \ldots]$	w_1
$\vdots$	$\vdots$	$\vdots$	$\vdots$
Q_m	$\mathcal{A}_m = \{A_m^1, A_m^2, \ldots\}$	$[\pi_m(A_m^1), \pi_m(A_m^2), \ldots]$	w_m
	Possible Prizes		
N/A	$X = \{x, x', x'', \ldots\}$	$[\pi_X(x), \pi_X(x'), \ldots]$	N/A

[a] Answers to different questions are not generally independent. Typically, the joint probability measure $\pi \neq \pi_1 \cdots \pi_m \cdot \pi_X$.

may think of beliefs as having intrinsic value, which is then amplified by these attention weights. Beliefs have positive (or negative) intrinsic value when a person likes (or dislikes) thinking about them, that is, when more attention enhances (or depresses) utility.

It is useful to distinguish two sources of a belief's intrinsic value: *valence* and *clarity*. Valence refers to the value of definitive answers to questions (see Brendl & Higgins, 1996). To illustrate the concept of valence, we return to the example of a professor's belief that she is a good (or bad) teacher as one with intrinsically positive (or, respectively, negative) valence. Smith, Bernheim, Camerer, and Rangel (2014) shows that valences (of beliefs about consumption) can be inferred from neural activity.

Clarity refers to preferences between degrees of certainty, independent of the answers one is certain of (see Kaplan, 1991). We assert that, *ceteris paribus,* people prefer to have greater clarity (i.e., less uncertainty or more definitive beliefs). The aversion that people feel toward uncertainty is reflected in neural responses in the anterior cingulate cortex, the insula and the amygdala (Hirsh & Inzlicht, 2008; Sarinopoulos et al., 2010).[13] It manifests in physiological responses as well. Subjects who know to expect an electric shock, but who are uncertain whether it will be mild or intense, show more fear—they sweat more profusely, and their hearts beat faster—than subjects who know for sure that an intense shock awaits (Arntz et al., 1992). The desire for clarity is consistent with an underlying drive for simplicity and sense-making (Chater & Loewenstein, 2015). When valence and clarity pull in opposite directions, it may be the case that people prefer a certain answer to a belief that dominates it on valence or that people prefer uncertainty when it leaves space for better answers.

A useful measure of the uncertainty about a particular question is the entropy of the probability distribution over answers (Shannon, 1948; see also Cabrales et al., 2013). The entropy of a (marginal) probability π_i is $H(\pi_i) = -\sum_{A_i \in \mathcal{A}_i} \pi_i(A_i) \log \pi_i(A_i)$ (with the convention that $0 \log 0 = 0$).[14] At one extreme with minimal entropy 0, there is a single answer known for sure; at the other extreme, for a question with finitely many answers, a uniform distribution maximizes entropy. Entropy, weighted by attention, satisfies Berlyne's (1957) criteria for a measure of the internal conflict or dissonance in one's cognitive state. We associate a psychological cost for beliefs with higher entropy as an instantiation of the desire for clarity.

A Specific Utility Function

To make precise predictions we consider a specific utility function incorporating a concern about material outcomes as well as a preference for beliefs characterized by high valence and clarity, taking the strength of this preference to be proportional to attention weights:

$$u(\pi, \mathbf{w}) = \sum_{x \in X} \pi_X(x) v_X(x) + \sum_{i=1}^{m} w_i \left(\sum_{A_i \in \mathcal{A}_i} \pi_i(A_i) v_i(A_i) - H(\pi_i) \right). \quad (1)$$

[13] Also, monkeys' dopamine neurons fire, indicating an intrinsic reward, when they have the opportunity to reduce uncertainty about the amount of water they will be given (Bromberg-Martin & Hikosaka, 2009).

[14] The base of the logarithm in the entropy formula is arbitrary and amounts to a normalization parameter.

The first term, the material component of the utility function, is the expected valuation of prizes. The second term, the belief-based component of the utility function, adds the value of beliefs weighted by attention, where the value of a belief consists of the expected valence of the answers minus its entropy.

We now describe properties (some quite strong and almost certainly not always satisfied) that characterize (and necessarily imply) this utility function (see Theorem 1 below).

Properties

The utility function in Equation (1) satisfies the following seven properties.[15]

Separability between questions. Additive separability of utility between questions means that a person can place a value on a belief about a given question without needing to consider beliefs about other questions.

P1. A utility function satisfies additive separability between questions if $u(\pi, \mathbf{w}) = u_X(\pi_X) + \sum_{i=1}^{m} u_i(\pi_i, w_i)$.[16]

Property (P1) may seem quite strong because we can imagine representations of sensible preferences that are not additively separable. For example, the value of a belief about whether a car on sale has a warranty intuitively could depend on the cost of the car in the first place (not to mention one's desire for a new car, one's estimation of the costs of car repairs, etc.). However, we may be able to represent these preferences as separable after all. We might suppose that these beliefs do have separable values but that they correlate with some other highly valued belief, perhaps about how good a deal one can get on the car. That is, while intuition tells us that the value of beliefs about different questions (e.g., "Does she like me?" and "Does she have a boyfriend?") is often interdependent, this dependence may be mediated by the existence of additional questions (e.g., "Will she go out with me?"), beliefs about which may be mutually dependent, but independently valued.

Expected valuation of prizes. Our theoretical framework assumes independence across cognitive states, giving us expected utility over cognitive states. Apart from the utility derived from beliefs, expected utility might extend to prizes as well.

P2. When the material component of the utility function is separable, it satisfies expected valuation of prizes if $u_X(\pi_X) = \sum_{x \in X} \pi_X(x) v_X(x)$.

Property (P2) implies that when a lottery is independent of beliefs about the world, its utility is simply the expected value of the prizes. However, to the extent that a gamble depends in part on some beliefs, the utility derived from these beliefs will also be relevant.

Monotonicity with respect to attention weights. Preferences satisfy the property of monotonicity with respect to attention weights if whenever increasing attention on a given belief enhances (or diminishes) utility, it will do so regardless of the absolute level of attention weight. At a psychological level, the interpretation of this monotonicity property is that when a belief is positive, more attention to it is always better, and when a belief is negative, more attention is always worse. In fact, the property provides a natural *definition* of whether a belief is positive or negative.

P3. Preferences satisfy monotonicty with respect to attention weights if for any $\mathbf{w}$, $\hat{\mathbf{w}}$, and $\hat{\hat{\mathbf{w}}} \in \mathbb{R}_+^m$ such that $w_i = \hat{w}_i = \hat{\hat{w}}_i$ for all i j and $\hat{\hat{w}}_j > \hat{w}_j > w_j$, we have $u(\pi, \hat{\mathbf{w}}) \geq u(\pi, \mathbf{w})$ if and only if $u(\pi, \hat{\hat{\mathbf{w}}}) \geq u(\pi, \hat{\mathbf{w}})$, with equality on one side implying equality on the other, for all $\pi \in \Delta(\alpha)$.

In the case that these inequalities hold strictly, we say that π_j, the belief about question Q_j, is a *positive belief*. If they hold as equalities, we say that π_j is a *neutral belief*. And, in the case that the inequalities hold in the reverse direction, then π_j is a *negative belief*.

Linearity with respect to attention weights. The next property describes how changing the attention on a belief impacts utility. For any given attention weight, the marginal utility of a change in belief depends on what those beliefs are and how much the individual values them. The property of linearity with respect to attention weights means that, in general, the mar-

[15] While our utility function violates Savage's (1954) sure-thing principle, it does satisfy a weaker "one-sided" version of it presented in the Appendix.

[16] A subset of questions $\tilde{\mathcal{Q}} \subset \mathcal{Q}$ can also be separable, in which case $u(\pi, \mathbf{w}) = \sum_{i:Q_i \in \tilde{\mathcal{Q}}} u_i(\pi_i, w_i) + u_{-\tilde{\mathcal{Q}}}(\pi_{-\tilde{\mathcal{Q}}}, \mathbf{w}_{-\tilde{\mathcal{Q}}})$ where $\pi_{-\tilde{\mathcal{Q}}}$ is the marginal distribution over answers to the remaining questions and prizes and the vector $w_{-\tilde{\mathcal{Q}}}$ contains the remaining components of $\mathbf{w}$.

ginal utility associated with such a change in belief (assuming the utility of this belief is separable) is proportional to the attention on that belief.

P4. When the utility of question Q_i is separable, linearity with respect to attention weights is satisfied if for any w_i and $\hat{w}_i \in \mathbb{R}_+$ and π_i' and $\pi_i'' \in \Delta(\mathcal{A}_i)$, we have

$$u_i(\pi_i', \hat{w}_i) - u_i(\pi_i'', \hat{w}_i) = \frac{\hat{w}_i}{w_i}\left(u_i(\pi_i', w_i) - u_i(\pi_i'', w_i)\right).$$

Property (P4) allows us, in the case of separable utility, to assign an intrinsic value v to beliefs such that $u_i(\pi_i', w_i) - u_i(\pi_i'', w_i) = w_i(v_i(\pi_i') - v_i(\pi_i''))$. We abuse notation by referring to the valence of answer A_i as $v_i(A_i)$, with it being defined here as the intrinsic value v_i of belief with certainty in A_i. We have taken the liberty of specifying a precise relationship between attention weights and utility as a convenient simplification; it should be noncontroversial because we do not claim to have a particular cardinal measure of attention weight.

Label independence. Intuitively, the value of a belief should depend on how an individual values the possible answers and on how probable each of these answers is, and these factors (controlling for attention weight of course) should be sufficient to determine the utility of any (uncertain) belief. In particular, the value of a belief should not depend on how the question or the answers are labeled.

P5. Label independence is satisfied if, when the utility of questions Q_i and Q_j are separable, a bijection $\tau : \mathcal{A}_i \to \mathcal{A}_j$, such that $v_i(A_i) = v_j(\tau(A_i))$ and $\pi_i(A_i) = \pi_j(\tau(A_i))$, implies that $v_i(\pi_i) = v_j(\pi_j)$.

Reduction of compound questions. The intuition behind the assumption of label independence also seems to suggest that the utility of a belief perhaps should not depend on the way the question giving rise to the belief is asked, that is, on whether a complicated question is broken up into pieces. We should recall, however, that the activation of a particular question directs attention to the belief about this question. Thus, in general, the utility of a belief will not be invariant to the question being asked. Still, it may be the case that utility remains invariant when a compound question is broken into parts as long as the attention on each part is weighted properly. If utility remains invariant upon setting attention weights on conditional questions to be proportional to the probabilities of the hypothetical conditions, then we say that the utility function satisfies the reduction of compound questions property. Figure 2 demonstrates the reduction of a compound question with appropriate attention weights on each subquestion.

P6. A separable utility function satisfies the reduction of compound questions property if whenever there is a partition ζ of the answers $\mathcal{A}_i$ (to question Q_i) into $\zeta = \{\mathcal{A}_{i_1}, \ldots, \mathcal{A}_{i_n}\}$ and a bijection $\tau : \zeta \to \mathcal{A}_j$ into the answers to some question Q_j such that for any $h \in [1, n]$ and any $A_i \in \mathcal{A}_{i_h}$,

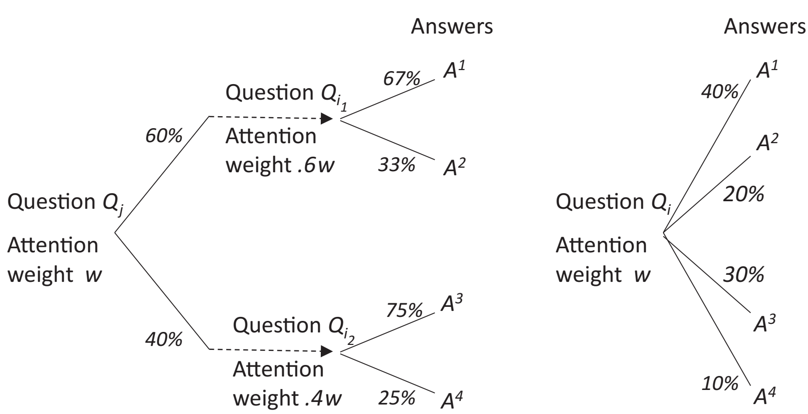

Figure 2. Decomposition of a compound question.

$$v_i(A_i) = v_j(\tau(\mathcal{A}_{i_h})) + v_{i_h}(A_i)$$

and

$$\pi_i(A_i) = \pi_j(\tau(\mathcal{A}_{i_h})) \cdot \pi_{i_h}(A_i),$$

It follows that

$$u_i(\pi_i, \omega) = u_j(\pi_j, \omega) + \sum_{h=1}^{n} u_{i_h}(\pi_{i_h}, \pi_j(\tau(\mathcal{A}_{i_h})) \cdot \omega).$$

Ruling out unlikely answers increases clarity. A final property operationalizes the preference for clarity. Controlling for the valence of one's beliefs, by considering situations in which one is indifferent between different possible answers to a question, there should be a universal aversion to being uncertain about the answer to an activated question. As a building block toward quantifying the uncertainty in a belief, we assert here that when an unlikely (and equally attractive) answer is ruled out, uncertainty decreases (and thus the utility of that uncertain belief increases).

P7. Ruling out unlikely answers increases clarity if, when the utility of question Q_i is separable and all answers to this question have the same valence, that is, $v_i(A_i) = v_i(A_i')$ for all A_i and $A_i' \in \mathcal{A}_i$, then for any $\boldsymbol{\pi}$ where without loss of generality $\pi_i(A_i^h)$ is weakly decreasing in h and for any π' such that $\pi_i'(A_i^h) \geq \pi_i(A_i^h)$ for all $h \in [1, \bar{h}]$ (with at least one inequality strict) and $\pi_i'(A_i^h) = 0$ for all $h > \bar{h}$, for some $\bar{h}$, we consequently have $v_i(\pi_i') > v_i(\pi_i)$.

Characterization of Our Utility Function

Theorem 1. If the properties P1–P7 are satisfied, then the utility function takes the form of Equation (1).

Proof. Linearity with respect to attention weights allows us to pull an attention weight on question Q_i outside of the utility $u_i(\pi_i, w_i) = w_i v_i(\pi_i)$ (using a neutral belief to calibrate v_i). A partition of $\mathcal{A}_i$ into singletons $\mathcal{A}_{i_h}$ such that $v_i(A_i) = v_{i_h}(A_i)$ allows us, by reduction of the compound question, to determine that the function $F(\pi_i) = v_i(\pi_i) - \sum_{A_i \in \mathcal{A}_i} \pi_i(A_i) v_i(A_i)$ does not depend on $v_i(A_i)$ for any $A_i \in \mathcal{A}_i$. Moreover, $-F(\cdot)$ satisfies Shannon's (1948) axioms (continuity, increasing in the number of equiprobable answers, and reduction of compound questions) characterizing the entropy function $H(\pi_i) = -\sum_{A_i \in \mathcal{A}_i} \pi_i(A_i) \log \pi_i(A_i)$.

Choices of Actions

People cannot just choose to put themselves in their most preferred cognitive state in $\mathcal{C}$. The decision variables in our model are actions, such as whether or not to acquire information or make a wager, which will influence beliefs and/or attention. Actions, in general, are operators on cognitive states that map to new cognitive states or to distributions over cognitive states.

Choosing Between Sequences of Actions

An initial choice is often made in the context of subsequent actions that will, or may, become available, as well as subsequent events beyond one's control that may also operate on one's cognitive state. For example, a college professor deciding whether or not to review her teaching ratings may subsequently have the option to enroll in a teacher improvement class. Similarly, a gambler deciding whether to wager on a football game may subsequently have the option to watch the game or just check the final score.[17]

A sequence of actions (and events) can be analyzed with the convention that an operator passes through a distribution over cognitive states.[18] Thus, we represent a sequence s of actions and events acting on a cognitive state $(\pi, \mathbf{w})$ as $s \cdot (\pi, \mathbf{w}) \in \Delta(\mathcal{C})$. Choice from a strategy set $\mathcal{S}$, that is, from a set of sequences of (possibly state-contingent) actions where early actions may reveal information that will inform later actions, is represented as utility maximization: A sequence $s^* \in \mathcal{S}$ may be chosen by a

[17] A person may be naive, that is, may not recognize that some subsequent action or event may take place, in which case the choice should be modeled without it.

[18] Analogous to the standard assumption in decision under risk, the model assumes reduction of compound distributions over cognitive states. This does not imply the traditional reduction of compound lotteries.

decision maker in the cognitive state (π, $\mathbf{w}$) if $s^* \in \operatorname{argmax}_{s \in \mathcal{S}} u(s \cdot (\pi, \mathbf{w}))$.[19]

Expected utility over cognitive states implies dynamic consistency for sequences of choices. For example, if the college professor intends to review her teaching ratings and enroll in the teacher improvement class if and only if less than half of the class liked her teaching, then after discovering, say, that everybody liked her teaching, she would not want to enroll in the class (or, conversely, if in this scenario she would want to enroll in the class, then she would have intended to do so as part of her initial plan).[20]

We can make use of dynamic consistency to break a complicated dynamic decision problem into pieces that are easier to analyze. We may begin by considering an initial choice, counting on subsequent choices to maximize utility in each contingency. We find it useful to define a utility function over cognitive states, contingent on the set of strategies that may subsequently be chosen:

$$U(\pi, \mathbf{w} \mid \mathcal{S}) = \max_{s \in \mathcal{S}} u(s \cdot (\pi, \mathbf{w})). \qquad (2)$$

In the example of the professor's teaching ratings, the initial choice is whether to review the teaching ratings, and the set of available subsequent actions is to enroll in the teacher improvement class or not to enroll in the class. Looking at the ratings resolves a lottery over cognitive states, each of which confers utility that is conditional on making the optimal choice of one of these subsequent actions. This dynamic decision problem is harder to grasp as one large decision. We have to define contingent plans of action: Enroll in the teacher improvement class if and only if the number of students who liked the teacher is in some range. The full strategy set then includes all such contingent plans of action (along with actually reviewing the teaching ratings) as well as not reviewing the teaching ratings and independently enrolling or not enrolling.[21]

Examples of Actions

We distinguish two kinds of actions: *informational* actions answer a question; *instrumental* actions affect the chances of receiving various prizes (outcomes).[22] For example, wagering on the color of a ball drawn from an urn is an instrumental action. Examining the contents of the urn is an informational action. Note that some actions will have both instrumental and informational effects. Examples include paying a fee for a property value appraisal or hiring a private eye.

A purely instrumental action acting on the prior cognitive state determines a particular new cognitive state. Typically, it preserves the prior judgment about the probability of each answer set and then specifies a new distribution over prizes conditional on each possible answer set. An instrumental action may also affect the importance of various questions (as formalized in the next section) and thereby influence the attention weights. For example, the decision to participate in a karaoke session will likely raise the attention weight on the question "Am I a good singer?"

Acquiring information also changes one's cognitive state. Ex ante, as one does not know which answer will be discovered, the prospect of acquiring information offers the decision maker a lottery over cognitive states. Upon learning answer A_i to question Q_i, one's probability measure over $\Delta(\alpha)$ changes from π^0 to $\pi^{A_i} = \pi^0(\cdot \mid A_i)$. We assume Bayesian updating here, which means that ex ante, before one

[19] If a sequence of actions produces a temporal profile of cognitive states, we can of course discount the utility of cognitive states reached in the future relative to those occurring in the present, but for simplicity we present our formal model with no time delay for subsequent actions.

[20] We assume expected utility over cognitive states, not over material outcomes. Thus, a person may, for example, plan to accept a compound lottery and then observe the final outcome, whereas if there were an opportunity to observe the outcome of the first stage and then reconsider the compound lottery, he may have rejected it (perhaps regardless of the outcome of the first stage). Observing the first stage is a consequential action in our model, and sequences of actions are absolutely not assumed to be commutative. Recognizing cognitive states as our objects of preference, accepting the compound lottery is not equivalent to accepting and observing the successive stages, so we have no violation of dynamic consistency here. See Machina (1989) for an illuminating discussion of consequentialism and dynamic consistency.

[21] If there are n students who have filled out teaching ratings, and thus $n + 1$ possible answers to how many students like the teacher, then the full strategy set includes $2^{n+1} + 2$ possible sequences of actions!

[22] A third kind of action, which we might call *meditative*, only focuses or relaxes attention. For example, dwelling on the contents of an urn (or removing it from sight) would be a meditative action.

knows what one will discover, an informational action determines a distribution over subjective judgments such that the expectation of this distribution equals the prior judgment. That is, by the law of total probability, $\sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i)\pi^{A_i} = \pi^0$. An informational action would decrease expected entropy because conditioning reduces entropy (see, e.g., Cover & Thomas, 1991, p. 27). New information generates surprise (as formalized in the next section), which changes the attention weights too. Given the prior attention weight vector $\mathbf{w}^0$, we let $\mathbf{w}^{A_i}$ denote the new attention weight vector immediately after learning A_i, resulting from surprise at this discovery. We make special note of a particular event that may follow after an informational action: adapting, that is, relaxing attention after getting used to new beliefs. We discuss how adaptation leads to decreased attention in the next section.

Determinants of Attention

To apply our framework in settings in which attention cannot be directly observed, we need to spell out how an action influences the attention paid to each activated question.[23] In this section we propose assumptions about the determinants of attention that describe how much people will think about various information gaps before and after taking some action.

Beginning with William James (1890), psychologists have drawn a distinction between automatic and voluntary control of attention. In our framework, in some situations, a person can choose to focus attention on (or ignore) a question, as when a reader of a mystery novel pauses to ponder who did it. In many familiar situations, however, attention is focused involuntarily, as when an anxious assistant professor wonders whether he will get tenure.

We formalize the concepts of *importance*, *salience*, and *surprise*, all of which, we assume, contribute to attention weight when attention is automatic. The importance γ_i of a question Q_i reflects the degree to which one's utility depends on the answer. Salience, distinctly, reflects the degree to which a particular context highlights the question. We denote the salience of question Q_i as $\sigma_i \in \mathbb{R}_+$. Finally, surprise is a factor that reflects the dependence of attention on the dynamics of information revelation, and specifically on the degree to which receiving new information changes one's beliefs. We denote the surprise associated with a revised belief about question Q_i as δ_i.

Assumption A1. We assume that the attention w_i on an activated question Q_i is strictly increasing in this question's importance γ_i, its salience σ_i, and the surprise δ_i associated with it.

Importance

The importance of a question reflects how much the answer matters to the person. That is, a question is important to the extent that one's utility depends on the answer. To be precise, we refer to the spread of the utilities associated with the different answers to a question.

Given a particular prior probability measure π^0 and a set $\mathcal{S}$ of sequences of actions available to the decision maker, the importance γ_i of question Q_i is a function (only) of the likelihood of possible answers and the utilities associated with these answers, captured as

$$\gamma_i = \phi\left(\left\langle \pi_i^0(A_i), U(\pi^{A_i}, \mathbf{w}^{A_i} \mid \mathcal{S}) \right\rangle_{A_i \in \mathrm{supp}\left(\pi_i^0\right)}\right)$$

where U is the utility function defined in Equation (2). We do not specify the precise form of this function ϕ, but instead impose the following conditions: We require that ϕ (i.e., importance) increases with mean-preserving spreads of the (subjective) distribution of utilities that would result from different answers to the question, and that it is invariant with respect to constant shifts of utility.

Thus, raising the stakes increases importance, as does increasing the range of potential outcomes (see Goldstein & Beattie, 1991). Consider, for example, the question "How much does (s)he like me?" (referring to a potential romantic partner). The importance of this question would increase either if the subject cared more about the partner's feelings—for example, if the subject desired a relationship—or if the range of possible answers to the question expanded—for example, if signals of strong interest alternated with signals raising doubt (see Givens, 1978, p. 349). On the other hand, if an

[23] We think of attention as observable (perhaps imperfectly) through some combination of eye tracking, brain scans, and self-reports, but choice experiments do not typically include such observations, which would be especially difficult to collect in the field.

answer is known with certainty, then by our definition nothing is at stake, so the underlying question is no longer important.

Admittedly, acquiring information should affect the importance of the questions being addressed, but this reconsideration of importance is not immediate. We model reconsideration of importance as part of adaptation to new belief rather than as an anticipated consequence of information acquisition.

Assumption A2. We assume that the importance of a question is updated if and when the decision maker adapts to new beliefs (rather than when the decision maker first updates these beliefs).

Our definition of importance is, by design, circular. Importance depends on utility, which in turn depends on the attention weight, but importance also contributes to attention weight. There is psychological realism to this circularity, which captures the dynamic processes giving rise to obsession: attention to a question raises its importance, and the elevated importance gives rise to intensified attention. If we assume that these processes unfold instantaneously, then importance (and, in turn, attention weight and utility) will be a fixed point of this composition of functions. In practice, we can make simple comparisons of importance without going to the trouble of specifying precise values.

Salience

The salience of a question depends on a variety of exogenous contextual factors. For example, a question could be salient if it has recently come up in conversation (i.e., it has been primed) or if other aspects of the environment remind an individual about it. Alternatively, a question could be more salient to an individual if the answer is, in principle, knowable, and even more so if other people around her know the answer but she does not. Comparison and contrast generally increase a question's salience (Duncan & Humphreys, 1989; Itti & Koch, 2001). Distractions, on the other hand, can decrease a question's salience (see Lavie, 2005).

Often a question may be salient despite being unimportant. For example, if someone asked an innocuous question such as "Is that plant a fern?" one might become curious about the answer despite not caring about whether it is "yes" or "no." It seems natural to think that some degree of salience is a necessary, and sufficient, condition for attention, while some degree of importance is not. One can, for example, often be curious about the answer to a question without preferring any particular answer over any other.

Assumption A3. We assume that a question Q_i is activated if and only if it has positive salience $\sigma_i > 0$.

Further, it seems natural to assume that salience and attention are positive complements; an increase in importance should produce a greater increase in attention for a more salient question, and vice versa.

Assumption A4. We assume that attention weight w_i has strictly increasing differences (i.e., a positive cross-partial derivative, if we assume differentiability) in (γ_i, σ_i).

Surprise

The third factor that we posit influences attention is the surprise one experiences upon acquiring new information. Surprise reflects the degree to which new information changes existing beliefs. We adopt Itti and Baldi's (2009) specification of surprise: When the answer to a particular question Q_j is learned, thereby contributing information about the answers to associated questions and causing their probabilities to be updated, the degree of surprise associated with a new belief about question Q_i can be defined as the Kullback-Leibler divergence of $\pi_i^{A_j}$ against the prior π_i^0,

$$\delta_i\left(\pi_i^{A_j} \,\Big\|\, \pi_i^0\right) = \sum_{A_i \in \mathcal{A}_i} \pi_i^{A_j}(A_i) \log \frac{\pi_i^{A_j}(A_i)}{\pi_i^0(A_i)}.$$

Surprise is positive with any new information, and is greatest when one learns the most unexpected answer with certainty. Itti and Baldi (2009; Baldi & Itti, 2010) show that surprise, specified this way, predicts the level of attention paid to information (see also Meyer et al., 1991). Mellers, Schwartz, Ho, and Ritov (1997) find that more surprising good (or bad) news is more elating (or disappointing) than less surprising news.

We return again to the example question "Do other people like me?" to illustrate surprise. If, having believed that she was generally well-liked, an individual were to discover that the comments about her were actually unfavorable, the discovery, necessitating a radical change in her belief, would be quite surprising (and would increase her attention to the question).

The feeling of surprise is not permanent.

Assumption A5. We assume that if and when the decision maker adapts to new beliefs, they are no longer surprising.

The Belief Resolution Effect

The impact of new information on attention is greatest when uncertainty about a question is resolved completely. Surprise generates an immediate spike in attention. However, with adaptation, surprise fades and the underlying question becomes unimportant because, with the answer known, there is no longer a range of possible answers. Taken together, these factors create a pattern of change in attention weight following the discovery of a definitive answer, what we call the *belief resolution effect*—when an answer is learned with certainty, there is an immediate boost in attention weight on it, but when the decision maker adapts, the question then receives less attention. It is as if the brain recognizes that because a question has been answered, it can move on to other questions that have yet to be addressed. Janis (1958) recognized the belief resolution effect when he observed that surgical patients getting information about their upcoming procedures initially worry more about the surgery but subsequently experience less anxiety. The belief resolution effect allows for hedonic adaptation to good or bad news (Smith et al., 2009; Wilson et al., 2005; Wilson & Gilbert, 2008). Kahneman and Thaler (2006, p. 230) explain:

> Withdrawal of attention is the main mechanism of adaptation to life changes such as becoming a paraplegic, becoming suddenly wealthy or getting married. Attention is normally associated with novelty. Thus, the newly paraplegic, lottery winner or newlywed is almost continuously aware of that state. But as the new state loses its novelty it ceases to be the exclusive focus of attention.

Often, however, people do not anticipate the belief resolution effect because they do not recognize that they will adapt (Wilson & Gilbert, 2005).

Information Acquisition and Avoidance

We can apply our utility function along with our assumptions about attention to decisions about information acquisition or avoidance. We develop our analysis in a companion article (Golman & Loewenstein, 2016) that uses our model to lay out three distinct motives for information acquisition or avoidance and derives the behavioral implications that result from the interplay of these motives. Here we summarize these implications and show with examples how our model allows for acquisition of noninstrumental information or avoidance of possibly useful information.

The desire for information, in our model, can be decomposed into three distinct motives: recognition of the instrumental value of the information, curiosity to fill the information gap(s), and motivated attention to think more or less about what could be discovered. The instrumental value of information arises from its impact on subsequent actions. As in the standard account of informational preferences, it is defined as the difference between the expected utility of subsequent actions conditional on having the information and the utility expected in the absence of the information. Curiosity arises from the expected reduction in uncertainty upon acquiring information. It is defined as the expected utility of revised beliefs, given prior levels of attention. The magnitude of curiosity depends on the attention devoted to each information gap that stands to be addressed. Motivated attention arises from the anticipated surprise (i.e., increase in attention to a question) upon acquiring information. It is defined as the expected utility from increased attention on whatever happens to be discovered, conditioning on all possible outcomes. Motivated attention is a motive to acquire information that's expected to be good and to avoid information that is expected to be bad.

Putting the three motives together, our model makes many predictions about when, and the degree to which, information will be sought or avoided. When anticipated answers are neutral or even potentially positive, information should be sought. The strength of the desire for this information should increase with the number of attention gaps that can be addressed, the attention paid to them, and the valence of the possible outcomes. However, when anticipated outcomes are sufficiently negative, information would be avoided.[24] This "ostrich effect" when anticipating bad outcomes is consistent with a growing body of empirical evidence (see, e.g., Eil & Rao, 2011; Falk & Zimmermann, 2016; Ganguly & Tasoff, 2016; Karlsson et al., 2009). For example, Ganguly and Tasoff (2016) find that willingness to pay to avoid testing for a herpes infection is greater for the more dreaded Type II infection than for Type I. According to our theory, information avoidance allows a person to escape from thinking about such negative beliefs. Falk and Zimmermann (2016) confirm this hypothesis, finding that information avoidance about impending electric shocks (clearly bad outcomes) is more prevalent when subjects can distract themselves by playing a game.

Information Acquisition: Curiosity

Judging by the demand for celebrity gossip blogs, many people are curious about the private lives of celebrities, be it whether a particular actress is pregnant again or whether a musician has relapsed into drug addiction or simply what clothing the celebrity was recently wearing in public. In many (yet not necessarily all) of these cases, an individual is curious despite having no use for the information and even not caring what the answer turns out to be. For example, a provocative headline may tempt a person to view a photo of Caitlyn Jenner's outfit to see whether she is wearing tight-fitting or loose-fitting jeans even if the person never intends to discuss the topic (or imitate the style) and is indifferent between these two possibilities. Our model is able to make sense of information acquisition due to curiosity because we explicitly model thoughts and feelings about the information gap.

Suppose "What is Caitlyn Jenner wearing?" is an activated (salient) question with possible answers tight jeans or loose jeans. The actual answer should not affect other beliefs or material outcomes at all. Consider the choice whether or not to find out the actual answer, with no other actions or choices subsequently available. We assume that knowing what she is wearing is a neutral belief, but the desire for clarity makes not knowing a negative belief. Learning the answer to the question reduces the entropy of this belief and thus raises utility. (Additional attention to her outfit, due to some surprise upon finding out, does not affect utility from neutral beliefs, and there is no instrumental value in the absence of subsequent actions.) Thus, the model yields a strict preference for finding out.

The concept of awareness embedded in the question–answer framework is crucial for the treatment of curiosity. It allows us to accommodate curiosity after reading the provocative headline along with a lack of curiosity before reading the headline. In both cases the person is uncertain about what Caitlyn Jenner is wearing, but only after reading the headline does the person become aware of this uncertainty (and hence become curious to resolve it).

Information Avoidance: The Ostrich Effect

Despite its instrumental value as well as the motive of curiosity, people sometimes avoid potentially useful information (see Golman, Hagmann, & Loewenstein, 2016 for a comprehensive review). For example, many at-risk individuals avoid getting tested for HIV despite the availability of life-saving treatment options (Thornton, 2008). Our model allows for information avoidance in this scenario (with bad to neutral outcomes anticipated) because getting a positive test result would lead to, and focus attention on, the negative belief that the individual has HIV.

Suppose "Do I have HIV?" is an activated question Q with possible answers yes or no. We might describe the relevant material outcomes in terms of quality-adjusted life years. The individual has prior belief π about the probability of having HIV and, conditional on this diagnosis, of the quality and length of life that can be expected. We denote the probability of having HIV as $p = \pi_Q(\text{yes})$. For simplicity of presentation, we consider the choice of whether or not to get tested as equivalent to a choice of whether or not to find out for sure if you have HIV. In reality, of course, diagnostic tests are not perfectly accurate, and a more careful analysis would consider the test results to be a distinct activated question that correlates strongly with

[24] The belief-resolution effect in our model also leads to a novel prediction: individuals with more foresight (and who discount the future less) should be less likely to exhibit the ostrich effect and more likely to acquire information despite anticipated bad news.

actually having HIV, but our simplified analysis more clearly demonstrates how the model works in application. Also for simplicity, we ignore other questions people might realistically have, such as how having HIV would affect their sex lives. The set of actions subsequently available includes an instrumental action s_M (taking medicine) or doing nothing. People do not typically take medicine to manage a potential HIV infection unless they get tested and actually have HIV, so we assume that $u(s_M \cdot (\pi, w)) < u(\pi, w)$ and $u(s_M \cdot (\pi^{no}, w^{no})) < u(\pi^{no}, w^{no})$, but $u(s_M \cdot (\pi^{yes}, w^{yes})) > u(\pi^{yes}, w^{yes})$.

We use the expected utility representation over cognitive states to express the expected utility after getting tested as $pu(s_M \cdot (\pi^{yes}, w^{yes}) + (1 - p)u(\pi^{no}, w^{no})$. Let us now adopt our proposed utility function specification from Equation (1). Let us set the material value of not having HIV (i.e., no reduction in the expected quality or length of one's life) to be 0 and denote the material value of the expected quality-adjusted life years with and without medicine as u_{XM} and u_{XH}, respectively, where $u_{XH} < u_{XM} < 0$. Let us also assume that having HIV has highly negative belief valence $v_H < 0$ and that not having HIV has neutral belief valence 0. We can now write the expected utility after getting tested as $p(u_{XM} + w^{yes} v_H)$. In comparison, the utility without getting tested can be written as:

$$pu_{XH} + w^0(pv_H + p\log p + (1 - p)\log(1 - p)).$$

The individual would get tested if and only if

$$p(u_{XM} - u_{XH}) + p(w^{yes} - w^0)v_H - w^0(p\log p + (1 - p)\log(1 - p)) \geq 0.$$

The first and last terms, corresponding to instrumental value and curiosity respectively, contribute positively. However, the middle term is negative because $w^{yes} > w^0$ due to surprise (and because $v_H < 0$). If thinking that you have HIV is sufficiently scary or unpleasant ($v_H << 0$), then this middle term dominates and causes the individual to avoid getting tested.

Risk and Ambiguity Preference

The previous section discusses how the model we have developed allows us to describe a desire to acquire or to avoid information that encompasses motives (namely, curiosity and motivated attention) that have been largely disregarded in the economics literature. We can apply this same model to an entirely new domain: preferences about wagers that depend on missing information. Risk and ambiguity preference are complex topics, and we develop these applications in depth in a companion article (Golman, Loewenstein, & Gurney, 2016). That article derives behavioral implications of the model in the domain of risk and ambiguity and reports an experimental test confirming our main new prediction. Here, we summarize those results and then apply the model to the Ellsberg paradox and to gambling for pleasure to show how the model proves useful in application.

Decision making under risk and under ambiguity both expose decision makers to information gaps. Imagine a choice between a gamble and a sure thing. Deciding to play the gamble naturally focuses attention on the question: what will be the outcome of the gamble? Of course, deciding to not play the gamble does not stop an individual from paying some attention to the same question (or, if not choosing the gamble means it will not be played out, the related question: what *would have been* the outcome of the gamble?) but playing the gamble makes the question more important, and that brings about an increase in the attention weight on the question. If the individual is aware of this effect, which it seems natural to assume, then whether it encourages risk taking or risk aversion will depend on whether thinking about the information gap is pleasurable or aversive. When thinking about the missing information is pleasurable, then the individual will be motivated to increase attention on the question, which entails betting on it. Conversely, when thinking about the missing information is aversive, the individual will prefer to not bet on it. This may help to explain why, for example, people generally prefer to bet on their home teams than on other teams, and especially when the other team is playing against their home team (Babad & Katz, 1991). A preference for betting on uncertainties that one likes thinking about shares much overlap with, but is distinguishable from, a preference for betting on uncertainties that one has expertise about (Heath & Tversky, 1991).

Decision making involving uncertainties that are ambiguous is similar to the case with known risks, but with an additional wrinkle: With ambiguity, there is an additional information gap. In a choice between a sure thing and an ambiguous gamble, for example, a second relevant question (in addition to the one above about the outcome of the gamble) is: What is the probability of winning with the ambiguous gamble? (And there may be additional relevant questions that could inform someone about this probability, so even a Bayesian capable of making subjective probability judgments would be exposed to these information gaps.) Again, betting on the ambiguous gamble makes these questions more important and thus will increase the attention weight on them. So, desire to play the gamble will be increasing with the degree to which thinking about the gamble is pleasurable. To the extent that abstract uncertainties are not pleasurable to think about, this model provides a novel account of standard demonstrations of ambiguity aversion, including those first generated by Ellsberg (1961) in his seminal article on the topic.

Ellsberg Two-Urn Paradox

In the Ellsberg two-urn paradox, subjects are presented with 2 urns. Urn 1 contains 100 red and black balls, but in an unknown ratio. Urn 2 has exactly 50 red and 50 black balls. Subjects must choose an urn to draw from, and bet on the color that will be drawn—they will receive a \$100 payoff if that color is drawn, and \$0 if the other color is drawn. Subjects must decide which they would rather bet on: (a) a red draw from Urn 1, or a black draw from Urn 1; (b) a red draw from Urn 2, or a black draw from Urn 2; (c) a red draw from Urn 1, or a red draw from Urn 2; and (d) a black draw from Urn 1, or a black draw from Urn 2. Intuition suggests that people will be indifferent between red and black in choices a and b, by the principle of insufficient reason, but will prefer Urn 2 to Urn 1 in choices c and d because this urn is less ambiguous. The axioms of subjective expected utility theory (Anscombe & Aumann, 1963; Savage, 1954), however, imply that a preference for an urn in choice c should imply a preference for the other urn in choice d, because the change in colors simply reverses winning and losing. Indeed, experimental evidence confirms the suspected violation of subjective expected utility theory (Becker & Brownson, 1964; MacCrimmon & Larsson, 1979).

According to our model, the desire for clarity, along with the desire to pay less attention to negative beliefs, would cause an individual to bet on the known urn rather than the ambiguous urn in the Ellsberg paradox. When a decision maker is presented with Ellsberg's choices, the following questions, among others, are activated:

- What is the composition of red and black balls in Urn 1?
- What is the composition of red and black balls in Urn 2?

Only the second question is known with certainty. Despite having no information from which to form an objective probability over answers to the first question, we assume the decision maker can form a subjective probability, and specifically that the decision maker is likely to assume a uniform distribution over possible compositions of Urn 1. Moreover, savvy decision makers will recognize that payoffs result from a compound lottery with Stage 1 determining the composition of the urn and Stage 2 determining the ball drawn from an urn with that composition, and they will reduce the compound lottery to form a belief that the prize will be won with probability.5. Nevertheless, a bet on a draw from Urn 1 makes the anticipated payoff contingent on the uncertain answer to the first question, whereas a bet on a draw from Urn 2 makes the anticipated payoff contingent on the certain answer to the second question.

We rely on three assumptions we have made: (a) attention weight on a question increases with the importance of that question; (b) increasing attention weight on a question with an unfavorable belief decreases utility; and (c) uncertain beliefs over answers to which one is indifferent are less favorable than certainty about one such answer. In this case, we assume no preference about the composition of an urn, independent of the eventual payoff, but there is of course the aforementioned preference for certainty. We assume that knowing the composition (of Urn 2), whatever it may be, is a neutral belief. The belief that Urn 1 has a uniform distribution over possible compositions, because of this uncertainty, is a negative belief. Thus, the decision maker prefers not to increase the attention weight on (the composition of) Urn 1 and can avoid doing so by choosing to bet on a draw from Urn 2 rather than a draw from Urn 1.

Recognizing feelings about information gaps allows us to explain the preference for betting on the known urn rather than on the unknown urn, even when the subjective probability judgment about the odds of winning a prize is the same for both urns. Our account relies on aversion to missing information rather than a distinction between objective and subjective probabilities. Also, the prediction of ambiguity aversion depends crucially on the maintained assumption that beliefs about the composition of an urn do not have positive valence. If the source of the ambiguity was an uncertainty that people enjoyed thinking about, the model would predict ambiguity-seeking behavior rather than ambiguity aversion.

Gambling for Pleasure

People sometimes voluntarily expose themselves to risk and ambiguity. For example, some football fans like to wager on the outcome of a football game. Our model accommodates gambling for pleasure in this scenario because placing the wager makes the outcome of the football game more important, and if a person enjoys thinking about football, this will increase his utility.

When a football fan has an opportunity to wager on a game, the question of who will win the game is activated.[25] Suppose, for simplicity, he believes each team has a 50% chance to win the game and is offered an even-money (1:1 odds) wager on the team of his choice. As a football fan, he enjoys thinking about who will win the game, so his (50/50) belief is a positive belief. Suppose the wager is modest enough that his valuation for the monetary outcomes is effectively linear. Then the wager would not affect the material component of his utility function because the expected monetary value of the wager is 0. Still, the wager can affect the belief-based component of his utility. If he has a rooting interest in the game, say if the Steelers winning has higher valence than the Browns winning, then the game is already somewhat important to the fan, but betting on the Steelers makes the outcome even more important (because the fan would gain even more if the Steelers win and lose even more if they lose), so he would prefer to bet on the Steelers. Additionally, if he has no rooting interest in the game, that is, equal valence for beliefs that either team will win, then the outcome is not important if he does not bet, but becomes important if he does. In this case, he would strictly prefer betting on either team to not betting at all.

Conclusion

In much of the economics literature, preferences about information have been viewed as derivative of risk preferences. We take a complementary perspective, considering thoughts and feelings about information gaps as primitive and viewing preferences about risk and ambiguity along with preferences about information as derivative of these thoughts and feelings.

Thoughts and feelings about information gaps underlie the acquisition of noninstrumental information as well as the avoidance of potentially useful information. People may obtain noninstrumental information purely to satisfy curiosity. Loewenstein (1994) proposed an information-gap account of curiosity, which provides insight about its situational determinants. There are many things that people do not know and that do not bother them, but awareness of specific pieces of missing information can prompt an unreasonably strong desire to fill these gaps. Our theory embraces the information gap concept and provides a new formal definition of an information gap. Our utility function assumes that people want to fill information gaps ceteris paribus (i.e., they desire clarity or dislike uncertainty), and this is a universal motive for information acquisition rather than avoidance. We identify this motive as curiosity. We hypothesize that information avoidance derives from a second motive, a desire to avoid increasing attention on a negative anticipated outcome. More generally, we suggest that individuals have an inclination to acquire (or avoid) information whenever they anticipate that what they discover will be pleasurable (or painful). Our fundamental assumption is that obtaining information tends to increase attention to it (as in Gabaix et al., 2006; Tasoff & Madarász, 2009) to the extent that it is surprising. This leads to the implication that people

[25] Many other related questions would realistically be activated as well, but we ignore them for simplicity of presentation. A more careful analysis would proceed similarly.

will seek information about questions they like thinking about and will avoid information about questions they do not like thinking about.

Research has shown that missing information has a profound impact on decision making under risk and ambiguity. For example, Ritov and Baron (1990) studied hypothetical decisions concerning whether to vaccinate a child, when the vaccine reduces the risk of the child dying from a disease but might itself be harmful. When the uncertainty was caused by salient missing information about the risks from vaccination—a child had a high risk of being harmed by the vaccine or no risk at all but it was impossible to find out which—subjects were more reluctant to vaccinate than in a situation in which all children faced a similar risk and there was no salient missing information. We suggest that ambiguity aversion in this scenario stems from the unpleasantness of thinking about this missing information (see also Frisch & Baron, 1988). This account of risk and ambiguity preference is conceptually different from, and has different testable implications from, the usual account of risk aversion involving loss aversion and the usual account of ambiguity aversion involving vague probabilities.[26] Our fundamental assumption relating risk and ambiguity preference to feelings about information gaps is that exposure to risk or ambiguity attracts attention to the underlying information gaps because it makes them more important. This leads to the implication that people will be averse to risk and ambiguity when they do not like thinking about the uncertainty and will seek risk and ambiguity when they like thinking about the uncertainty.

Preferences about information acquisition and about exposure to risk or ambiguity are often inextricably linked, and our theory makes predictions about their relationship. We can identify questions for which desire for information to answer the question should be positively or negatively correlated with desire to bet on what the answer will be.[27] Figure 3 displays our theoretical predictions as the valence of the answers to a question changes, in the special case that all answers are considered to have equal valence and equal probability and this belief is independent of other beliefs. When all answers have sufficiently negative valence, we would predict information avoidance and risk or ambiguity aversion; when all answers have sufficiently positive valence, we would predict information acquisition and risk or ambiguity seeking. When all answers have neutral valence, however, we would predict a negative correlation, that is, information acquisition and risk- or ambiguity aversion. Our theory makes the strong prediction that people will not exhibit information avoidance and risk or ambiguity seeking for the same question.

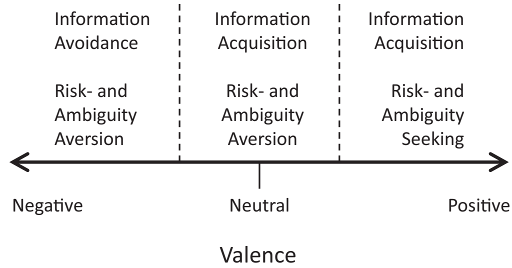

Figure 3. Theoretical predictions of informational preference and risk and ambiguity preference arising from thoughts and feelings about an information gap for which all answers have equal valence and equal probability, independent of other beliefs.

Additionally, relating informational preferences and risk and ambiguity preferences to thoughts and feelings about specific uncertainties (i.e., information gaps) provides coherence to the mixed findings about preferences about the timing of resolution of uncertainty (for gambles that subjects are already exposed to). Different studies have found varying idiosyncratic choices between early and late resolution as well as between gradual or one-shot resolution (Ahlbrecht & Weber, 1997; Falk & Zimmermann, 2016; Kocher, Krawczyk, & van Winden, 2014; Zimmermann, 2014). Our perspective suggests that people would like gradual resolution for uncertainties they enjoy thinking about (to savor them) and one-shot resolution for unpleasant uncertainties (to get them over with). Indeed, Kocher, Krawczyk, and van Winden (2014), using lottery jackpot drawings, find fairly frequent choices for gradual resolution whereas Falk and Zimmermann (2016), using electric shocks as outcomes, find that a majority of subjects opted for one-shot resolution. Also con-

[26] For example, low-stakes risk aversion (Rabin, 2000) could be attributed to the discomfort of thinking about uncertainties.

[27] We make no predictions about correlations between desire for information to answer one question and desire to bet on the answer to another, unrelated question.

sistent with our basic premise, anticipated feelings about waiting for a gamble to be resolved have been found to correlate with willingness to take the gamble in the first place (Loewenstein et al., 2001; Lovallo & Kahneman, 2000; van Winden et al., 2011).[28]

In this article we formally define information gaps and describe people's thoughts and feelings about them. We outline our theory's application to preferences for information acquisition or avoidance and preferences about risk and ambiguity, and we show with examples the usefulness of our theory for reconciling behavioral anomalies. We refer the reader to companion articles (Golman & Loewenstein, 2016; Golman, Loewenstein, & Gurney, 2016) that derive more general results and give each of these applications the attention they deserve.

[28] According to our model, risk and ambiguity preference can be affected by plans to observe the resolution of uncertainty. If beliefs about good or bad outcomes of a lottery (or gamble) not only correlate with the material component of utility but also have valence, then surprise will accentuate the hedonic impact of observing the outcome. Because surprise is larger for unexpected outcomes, unlikely events will have disproportionate impacts on expected utility. Much like overweighting of small probabilities in the original formulation of prospect theory (Kahneman & Tversky, 1979), this would lead to risk (and ambiguity) seeking for positively skewed lotteries (and gambles) and event-splitting violations of dominance (see Birnbaum, 2004).

References

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

Abelson, R. (1986). Beliefs are like possessions. *Journal for the Theory of Social Behavior, 16,* 223–250.

Ahlbrecht, M., & Weber, M. (1997). Preference for gradual resolution of uncertainty. *Theory and Decision, 43,* 167–185.

Andries, M., & Haddad, V. (2015). *Information aversion*. Paper presented at the Bounded Rationality in Choice Conference, New York, NY.

Anscombe, F., & Aumann, R. (1963). A definition of subjective probability. *Annals of Mathematical Statistics, 34,* 199–205.

Arntz, A., van Eck, M., & de Jong, P. (1992). Unpredictable sudden increases in intensity of pain and acquired fear. *Journal of Psychophysiology, 6,* 54–64.

Asch, D., Patton, J., & Hershey, J. (1990). Knowing for the sake of knowing: The value of prognostic information. *Medical Decision Making, 10,* 47–57.

Babad, E., & Katz, Y. (1991). Wishful thinking: Against all odds. *Journal of Applied Social Psychology, 21,* 1921–1938.

Baldi, P., & Itti, L. (2010). Of bits and wows: A Bayesian theory of surprise with applications to attention. *Neural Networks, 23,* 649–666.

Becker, S., & Brownson, F. (1964). What price ambiguity? Or the role of ambiguity in decision-making. *Journal of Political Economy, 72,* 62–73.

Benabou, R., & Tirole, J. (2002). Self-confidence and personal motivation. *Quarterly Journal of Economics, 117,* 871–915.

Berlyne, D. (1954). An experimental study of human curiosity. *British Journal of Psychology, 45,* 256–265.

Berlyne, D. E. (1957). Uncertainty and conflict: A point of contact between information-theory and behavior-theory concepts. *The Psychological Review, 64,* 329–339.

Berlyne, D. (1960). *Conflict, arousal, and curiosity*. New York, NY: McGraw-Hill.

Bernheim, B. D., & Rangel, A. (2009). Beyond revealed preference: Choice-theoretic foundations for behavioral welfare economics. *The Quarterly Journal of Economics, 124,* 51–104.

Birnbaum, M. (2004). Tests of rank-dependent utility and cumulative prospect theory in gambles represented by natural frequencies: Effects of format, event framing, and branch splitting. *Organizational Behavior and Human Decision Processes, 95,* 40–65.

Brendl, C., & Higgins, E. (1996). Principles of judging valence: What makes events positive or negative? In M. P. Zanna (Ed.), *Advances in experimental social psychology* (pp. 95–160). New York, NY: Academic Press.

Bromberg-Martin, E., & Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron, 63,* 119–126.

Cabrales, A., Gossner, O., & Serrano, R. (2013). Entropy and the Value of Information for Investors. *American Economic Review, 103,* 360–377.

Caplin, A., & Leahy, J. (2001). Psychological expected utility theory and anticipatory feelings. *Quarterly Journal of Economics, 116,* 55–79.

Caplin, A., & Leahy, J. (2004). The supply of information by a concerned expert. *Economic Journal, 114,* 487–505.

Chater, N. (2015). Can cognitive science create a cognitive economics? *Cognition, 135,* 52–55.

Chater, N., & Loewenstein, G. (2015). The under-appreciated drive for sense-making. *Journal of Economic Behavior and Organization, 126,* 137–154.

Cover, T., & Thomas, J. (1991). *Elements of information theory*. New York, NY: Wiley.

Dillenberger, D. (2010). Preferences for one-shot resolution of uncertainty and allais-type behavior. *Econometrica, 78,* 1973–2004.

Duncan, J., & Humphreys, G. W. (1989). Visual search and stimulus similarity. *Psychological Review, 96,* 433–458.

Dwyer, L., Shepperd, J., & Stock, M. (2015). Predicting avoidance of skin damage feedback among college students. *Annals of Behavioral Medicine, 49,* 685–695.

Eil, D., & Rao, J. (2011). The good news-bad news effect: Asymmetric processing of objective information about yourself. *American Economic Journal: Microeconomics, 3,* 114–138.

Eliaz, K., & Spiegler, R. (2006). Can anticipatory feelings explain anomalous choices of information sources? *Games and Economic Behavior, 56,* 87–104.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *Quarterly Journal of Economics, 75,* 643–699.

Falk, A., & Zimmermann, F. (2016). Beliefs and utility: Experimental evidence on preferences for information. (IZA Discussion Paper No. 10172).

Fantino, E., & Silberberg, A. (2010). Revisiting the role of bad news in maintaining human observing behavior. *Journal of the Experimental Analysis of Behavior, 93,* 157–170.

Fehr, E., & Rangel, A. (2011). Neuroeconomic foundations of economic choice—Recent advances. *Journal of Economic Perspectives, 25,* 3–30.

Fox, C., & Tversky, A. (1995). Ambiguity aversion and comparative ignorance. *Quarterly Journal of Economics, 110,* 585–603.

Frisch, D., & Baron, J. (1988). Ambiguity and rationality. *Journal of Behavioral Decision Making, 1,* 149–157.

Gabaix, X., Laibson, D., Moloche, G., & Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *American Economic Review, 96,* 1043–1068.

Ganguly, A., & Tasoff, J. (2016). Fantasy and dread: The demand for information and the consumption utility of the future. *Management Science*.

Geanakoplos, J., Pearce, D., & Stacchetti, E. (1989). Psychological games and sequential rationality. *Games and Economic Behavior, 1,* 60–79.

Givens, D. (1978). The nonverbal basis of attraction: Flirtation, courtship, and seduction. *Psychiatry: Interpersonal and Biological Processes, 41,* 346–359.

Gneezy, U., & Potters, J. (1997). An experiment on risk taking and evaluation periods. *Quarterly Journal of Economics, 112,* 631–645.

Goldstein, W., & Beattie, J. (1991). Judgments of relative importance in decision making: The importance of interpretation and the interpretation of importance. In D. Brown & J. E. K. Smith (Eds.), *Frontiers of mathematical psychology* (pp. 110–137). New York, NY: Springer.

Golman, R., Hagmann, D., & Loewenstein, G. (2016). Information avoidance. *Journal of Economic Literature*. Manuscript in preparation.

Golman, R., & Loewenstein, G. (2016). *The demand for, and avoidance of, information*. [Working Paper].

Golman, R., Loewenstein, G., & Gurney, N. (2016). *Information gaps for risk and ambiguity*. [Working Paper].

Grant, S., Kajii, A., & Polak, B. (1998). Intrinsic preference for information. *Journal of Economic Theory, 83,* 233–259.

Gul, F. (1991). A theory of disappointment aversion. *Econometrica, 59,* 667–686.

Hammond, P. (1988). Consequentialist foundations for expected utility. *Theory and Decision, 25,* 25–78.

Heath, C., & Tversky, A. (1991). Preference and belief: Ambiguity and competence in choice under uncertainty. *Journal of Risk and Uncertainty, 4,* 5–28.

Hirsh, J., & Inzlicht, M. (2008). The devil you know: Neuroticism predicts neural response to uncertainty. *Psychological Science, 19,* 962–967.

Hirshleifer, J., & Riley, J. (1979). The analytics of uncertainty and information—An expository survey. *Journal of Economic Literature, 17,* 1375–1421.

Hsee, C., & Ruan, B. (2016). The Pandora effect: The power and peril of curiosity. *Psychological Science, 27,* 659–666.

Hsu, M., Bhatt, M., Adolphs, R., Tranel, D., & Camerer, C. (2005). Neural systems responding to degrees of uncertainty in human decision-making. *Science, 310,* 1680–1683.

Itti, L., & Baldi, P. (2009). Bayesian surprise attracts human attention. *Vision Research, 49,* 1295–1306.

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience, 2,* 194–203.

James, W. (1890). *Principles of psychology*. New York, NY: Holt.

Janis, I. (1958). *Psychological stress: Psychoanalytic and behavioral studies of surgical patients*. New York, NY: Wiley.

Kahneman, D., & Thaler, R. (2006). Anomalies: Utility maximization and experienced utility. *Journal of Economic Perspectives, 20,* 221–234.

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

Kang, M. J., Hsu, M., Krajbich, I., Loewenstein, G., McClure, S., Wang, J., & Camerer, C. (2009). The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychological Science, 20,* 963–973.

Kaplan, S. (1991). Beyond rationality: Clarity-based decision making. In T. Gärling & G. W. Evans (Eds.), *Environment, cognition and action: An integrative multidisciplinary approach* (pp. 171–187). New York, NY: Oxford University Press.

Karlsson, N., Loewenstein, G., & Seppi, D. (2009). The ostrich effect: Selective attention to information. *Journal of Risk and Uncertainty, 38,* 95–115.

Karni, E., & Viero, M.-L. (2013). "Reverse Bayesianism:" A choice-based theory of growing awareness. *American Economic Review, 103,* 2790–2810.

Kimball, M. (2015). *Cognitive economics*. NBER Working Papers 20834.

Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica, 73,* 1849–1892.

Kocher, M., Krawczyk, M., & van Winden, F. (2014). Let me dream on! Anticipatory emotions and preference for timing in lotteries. *Journal of Economic Behavior & Organization, 98,* 29–40.

Köszegi, B. (2003). Health anxiety and patient behavior. *Journal of Health Economics, 22,* 1073–1084.

Köszegi, B. (2006). Ego utility, overconfidence, and task choice. *Journal of the European Economic Association, 4,* 673–707.

Köszegi, B. (2010). Utility from anticipation and personal equilibrium. *Economic Theory, 44,* 415–444.

Kreps, D., & Porteus, E. (1978). Temporal resolution of uncertainty and dynamic choice theory. *Econometrica, 46,* 185–200.

Kruger, J., & Evans, M. (2009). The paradox of alypius and the pursuit of unwanted information. *Journal of Experimental Social Psychology, 45,* 1173–1179.

Lavie, N. (2005). Distracted and confused?: Selective attention under load. *Trends in Cognitive Sciences, 9,* 75–82.

Lieberman, D., Cathro, J., Nichol, K., & Watson, E. (1997). The role of s- in human observing behavior: Bad news is sometimes better than no news. *Learning and Motivation, 28,* 20–42.

Litman, J., Hutchins, T., & Russon, R. (2005). Epistemic curiosity, feeling-of-knowing, and exploratory behaviour. *Cognition and Emotion, 19,* 559–582.

Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *The Economic Journal, 97,* 666–684.

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological Bulletin, 116,* 75–98.

Loewenstein, G., Weber, E., Hsee, C., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin, 127,* 267–286.

Lovallo, D., & Kahneman, D. (2000). Living with uncertainty: Attractiveness and resolution timing. *Journal of Behavioral Decision Making, 13,* 179–190.

MacCrimmon, K., & Larsson, S. (1979). Utility theory: Axioms versus "paradoxes". In M. Allais & O. Hagen (Eds.), *Expected utility hypotheses and the allais paradox* (pp. 333–409). Dordrecht, Holland: D. Reidel Publishing Company.

Machina, M. (1989). Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature, 27,* 1622–1668.

Marvin, C., & Shohamy, D. (2016). Curiosity and reward: Valence predicts choice and information prediction errors enhance learning. *Journal of Experimental Psychology: General, 145,* 266–272.

Mellers, B. A., Schwartz, A., Ho, K., & Ritov, I. (1997). Decision affect theory: Emotional reactions to the outcomes of risky options. *Psychological Science, 8,* 423–429.

Menon, S., & Soman, D. (2002). Managing the power of curiosity for effective web advertising strategies. *Journal of Advertising, 31*(3), 1–14.

Meyer, W., Niepel, M., Rudolph, U., & Schutzwohl, A. (1991). An experimental analysis of surprise. *Cognition & Emotion, 5,* 295–311.

Nielsen, L. (1984). Unbounded expected utility and continuity. *Mathematical Social Sciences, 8,* 201–216.

Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica, 68,* 1281–1292.

Ritov, I., & Baron, J. (1990). Reluctance to vaccinate: Omission bias and ambiguity. *Journal of Behavioral Decision Making, 3,* 263–277.

Sarinopoulos, I., Grupe, D., Mackiewicz, K., Herrington, J., Lor, M., Steege, E., & Nitschke, J. (2010). Uncertainty during anticipation modulates neural responses to aversion in human insula and amygdala. *Cerebral Cortex, 20,* 929–940.

Savage, L. (1954). *The Foundations of Statistics*. New York, NY: Wiley.

Schelling, T. (1987). The mind as a consuming organ. In J. Elster (Ed.), *The multiple self* (pp. 177–196). Cambridge, UK: Cambridge University Press.

Schweizer, N., & Szech, N. (2014). *Optimal revelation of life-changing information*. [Working Paper].

Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal, 27,* 623–656.

Smith, A., Bernheim, B. D., Camerer, C., & Rangel, A. (2014). Neural activity reveals preferences without choices. *American Economic Journal Microeconomics, 6*(2), 1–36.

Smith, D. M., Loewenstein, G., Jankovich, A., & Ubel, P. A. (2009). Happily hopeless: Adaptation to a permanent, but not to a temporary, disability. *Health Psychology, 28,* 787–791.

Stigler, G. (1961). The economics of information. *Journal of Political Economy, 69,* 213–225.

Tasoff, J., & Madarász, K. (2009). *A model of attention and anticipation*. [Working Paper].

Thornton, R. (2008). The demand for, and impact of, learning HIV status. *American Economic Review, 98,* 1829–1863.

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

van Dijk, E., & Zeelenberg, M. (2007). When curiosity killed regret: Avoiding or seeking the unknown in decision-making under uncertainty. *Journal of Experimental Social Psychology, 43,* 656–662.

van Winden, F., Krawczyk, M., & Hopfensitz, A. (2011). Investment, resolution of risk, and the role of affect. *Journal of Economic Psychology, 32,* 918–939.

von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton, NJ: Princeton University Press.

Wakker, P. (1988). Nonexpected utility as aversion of information. *Journal of Behavioral Decision Making, 1,* 169–175.

Wilson, T., Centerbar, D., Kermer, D., & Gilbert, D. (2005). The pleasures of uncertainty: Prolonging positive moods in ways people do not anticipate. *Journal of Personality and Social Psychology, 88,* 5–21.

Wilson, T., & Gilbert, D. (2005). Affective forecasting: Knowing what to want. *Current Directions in Psychological Science, 14,* 131–134.

Wilson, T., & Gilbert, D. (2008). Explaining away: A model of affective adaptation. *Perspectives on Psychological Science, 3,* 370–386.

Yariv, L. (2001). *Believe and let believe: Axiomatic foundations for belief dependent utility functionals*. Cowles Foundation Discussion Paper No. 1344.

Zimmermann, F. (2014). Clumped or piecewise? Evidence on preferences for information. *Management Science, 61,* 740–753.

Appendix

The One-Sided Sure-Thing Principle

The one-sided sure-thing principle asserts that people always prefer a certain answer to uncertainty among answers that all have valences no better than the certain answer (holding attention weight constant):

For any $\pi \in \Delta(\alpha)$, let $\mathrm{supp}(\pi) \subseteq \alpha$ denote the support of π. If for all $\mathbf{A} \times x \in \mathrm{supp}(\pi)$ we have $u(\pi', \mathbf{w}) \geq u(\pi^{\mathbf{A}\times x}, \mathbf{w})$, then $u(\pi', \mathbf{w}) \geq u(\pi, \mathbf{w})$, with the latter inequality strict whenever there exist $\mathbf{A}' \times x'$ and $\mathbf{A}'' \times x'' \in \mathrm{supp}(\pi)$ such that $\mathbf{A}' \neq \mathbf{A}''$.

Received January 25, 2016
Revision received June 21, 2016
Accepted August 12, 2016 ■