Curiosity, Information Gaps, and the Utility of Knowledge

Russell Golman and George Loewenstein [*][†]

August 4, 2014

Abstract

We propose an integrated theoretical framework that captures the diverse motives driving the preference to obtain or avoid information. Beyond the conventional desire for information as an input to decision making, people are driven by curiosity, which is a desire for knowledge for its own sake, even in the absence of material benefits, and people are additionally motivated to seek out information about issues they like thinking about and avoid information about issues they do not like thinking about (an "ostrich effect"). The standard economic framework is enriched with the insights that knowledge has valence, that ceteris paribus people want to fill in information gaps, and that, beyond contributing to knowledge, information affects the focus of attention.

KEYWORDS: curiosity, information, information gap, motivated attention, ostrich effect

JEL classification codes: D81, D83

[*]Department of Social and Decision Sciences, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA. E-mail: rgolman@andrew.cmu.edu; gl20@andrew.cmu.edu

[†]We thank Botond Köszegi, David Laibson, Erik Eyster, Juan Sebastian Lleras, Alex Imas, Leeann Fu, and Donald Hantula for helpful comments.

1 Introduction

In a seminal paper titled "The Mind as a Consuming Organ," Thomas Schelling (1987) pointed out that much consumption is not of the material sort, but takes place largely "in the mind." Standard accounts of utility maximization tend to treat information and consumption separately and radically differently. Specifically, the standard economic theory of information (Stigler, 1961) assumes that people seek out information because, and only to the extent that, it enables them to make superior decisions that raise their expected utility. However, research in psychology, decision theory, and most recently economics, has identified a number of other motives underlying the demand for information, from the powerful force of curiosity (Loewenstein, 1994) to the pleasures of knowledge and insight (Karlsson et al., 2004) apart from any material benefits they might confer. The purpose of this paper is to enumerate these different motives and show that they can be assimilated into a unified theoretical framework in which information not only informs decision making but also directly impacts utility by changing beliefs and redirecting the focus of attention. The theory we develop has a variety of implications that distinguish it from the standard economic theory of the demand for information.

A major divergence between our theory and standard economic theory is that our theory can account for the many situations in real life in which people actively *resist* acquiring information. For example, people often choose to not obtain medical tests, even when the test is costless (e.g., simply checking a box when giving a blood sample) and even when the results of the test could provide valuable information for decision making (e.g., whether to obtain treatment) (Lyter et al., 1987; Lerman et al., 1996; Lerman et al., 1999; Caplin and Eliaz, 2003; Köszegi, 2003; Thornton, 2008; Oster et al., 2013; see also Sweeny et al., 2010). The standard economic account of the value of information predicts that (outside of strategic situations) valid information will never be valued negatively since, at worst, it can be ignored, i.e., not taken into account in decision making. If beliefs enter into the utility function, however, it follows that some information might not only not be valued, but be valued negatively.

Previous treatments of information avoidance in economics have generally derived information-avoidance from risk preferences (e.g., Kreps and Porteus, 1978; Wakker, 1988; Grant et al., 1998; Dillenberger, 2010) or disappointment aversion (e.g., Köszegi, 2006; Karlsson et al., 2009; Köszegi, 2010).[1] In contrast, our model hypothesizes that information avoidance derives from the desire to avoid increasing attention on a negative anticipated outcome. While we also hypothesize that awareness of uncertainty is aversive (which, we argue, is the underlying cause of curiosity), the desire to distract attention from potentially unpleasant beliefs can, in some situations, trump curiosity and lead to information avoidance.

The different motives driving the desire for, or desire to avoid, information that our theory encompasses can be depicted in a two-by-two matrix (Table 1; labeled with letters indicating the order in which we discuss them). The left-hand column encompasses desire for (or against) information that is extrinsic – i.e., connected to decision making. The right-hand column encompasses situations in which information may be sought or avoided for itself – i.e., independently of any expected impact on decision making. The distinction between the rows is somewhat more subtle. The top row, labeled "valence" (Brendl and Higgins,

[1] Information avoidance has also been rationalized as a self-control mechanism (see Carillo and Mariotti, 2000; Bernheim and Thomadsen, 2005; Dana et al., 2007), much like how information might be avoided in strategic settings to influence other players.

Motivation	Connected to Decision Making Yes (extrinsic)	 No (intrinsic)
Valence	A) instrumental value (standard economics)	B) motivated attention to (and avoidance of) information
Clarity (addressing an 'information-gap')	D) risk- and ambiguity-preference	C) curiosity

Table 1: Two-dimensional representation of the demand for (and against) information.

1996), recognizes that updating one's intended course of action, or even just one's beliefs when information is acquired, may affect utility in a directional fashion. The bottom row, labeled "clarity" (Kaplan, 1991), recognizes a natural desire to fill an 'information gap' (Loewenstein, 1994) – an awareness of the absence of potentially useful or interesting information. Uncertainty about such missing information is aversive, even if it doesn't matter what one might find.

The top left-hand box represents the customary economic account of the demand for information, in which information is valued because, and only to the extent that, it results in decisions with superior expected outcomes.

The top right-hand box captures situations in which individuals seek or avoid information because they anticipate that what they discover will be pleasurable or painful - i.e., will lead to an increase or decrease in utility. From a Bayesian perspective, it might seem strange that a decision maker would expect that obtaining information, which by its very nature is not known, would have a non-zero expected impact on utility. The assumption that good news is pleasurable and bad news painful is not sufficient, on its own, to produce an ex ante preference for gathering information about events that are likely to be good and avoiding information about events that are likely to be bad, because ex-ante beliefs about such events are already good or bad respectively (Eliaz and Spiegler, 2006). Nevertheless, there can be a big difference between discovering something for sure and simply considering it a likely possibility. Our additional assumption is that obtaining news tends to increase attention to it (as in Gabaix et al., 2006; Tasoff and Madarász, 2009), which leads to the implication that people will seek information about questions they like thinking about and will avoid information about questions they do not like thinking about.

The bottom right-hand cell designates the phenomenon of curiosity, which is in the right-hand column because, by definition, it refers to the desire for information for its own sake – i.e., specifically *not* for its ability to improve decision making. 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. Our treatment of curiosity adopts an information gap account proposed by the second author (Loewenstein, 1994) that has found its way into the psychology, but not yet the economics, literature.

The bottom left-hand cell refers to cases in which one experiences an information gap that *is* relevant to a decision one faces. This situation is similar to the standard case (top left cell) except that the information is not obtainable. Research has shown that missing information has a profound impact on decision making. 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. In a companion paper (Golman and Loewenstein, 2014) we argue that the information-gap concept developed here underlies an alternative account of risk and ambiguity aversion (and seeking) that 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. In the current paper, in Section 5, we only outline our argument that salient information gaps can either increase or decrease preference for uncertain gambles depending on whether it is painful or pleasurable to think about the information one is missing.

Although all four of the motives represented in Table 1 have been addressed in the economics and/or psychology literature, they have generally been examined separately, and have never, to the best of our knowledge, been assimilated into a single integrative framework. Our theory captures a wide range of empirical phenomena and generates not only predictions that are consistent with well-documented patterns of behavior, but also a number of novel predictions that are testable, but as yet untested.

Our approach builds on the insights of Caplin and Leahy (2001). Caplin and Leahy recognize that anticipatory feelings about prizes that might be received in the future can affect utility. We follow them (and Köszegi (2010) as well) 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. In doing so, we incorporate Tasoff and Madarász's (2009) insight that information stimulates attention and thus complements anticipatory feelings. Kreps and Porteus (1978) present a model capturing preferences for early or late resolution of uncertainty, and Dillenberger (2010) captures preferences for one-shot or sequential resolution of uncertainty; this line of research thus deals with when, but not whether, an individual prefers to acquire information. Our model focuses just on the latter issue, but with it one could address the timing of uncertainty resolution by making additional assumptions about time preference.

Benabou and Tirole (2002) also allow people to care about their beliefs. In their model, beliefs about the self (i.e., self-confidence or self-esteem) affect motivation to complete a task that requires will power, and thereby beliefs play into the utility function. Köszegi's (2006) model has a similar flavor and allows beliefs about the self to directly affect utility. These models predict information avoidance when people have high self-confidence and don't want new information to taint this view. Our model, on the other hand, predicts information avoidance in different circumstances, specifically when people hold more negative beliefs. In Section 4.4 we describe conditions in which information avoidance is more likely to occur.

Brunnermeier and Parker (2005) advance a model in which people choose their beliefs. In Oster et al.'s (2013) application of Brunnermeier and Parker's model to testing for Huntington's Disease, they assume that people avoid getting tested so they can remain optimistic. Our model does not address optimism or pessimism, and assumes that ex-ante beliefs are the expectation of possible ex-post beliefs, in accordance with Bayesian updating. In our model people avoid obtaining information so as to not increase attention to something that is uncomfortable to think about, not to avoid updating beliefs. In Section 4.5 we provide an alternative account of test-avoidance based on people managing attention rather than maintaining optimism.

Benabou (2013) explores spillovers related to denial of bad news or information avoidance, identifying contexts in which such behaviors are likely to spread throughout a group of people. This work demonstrates that preferences about information, of the kind we consider here, have meaningful consequences for organizations and markets.[2]

We rely on a reduced form model of knowledge and awareness to describe information gaps – and the desire to fill them or ignore them – in order to avoid the complications of working with information partitions in a state-space model of knowledge (as in Aumann, 1976). The standard partitional state-space framework permits a distinction between two states of affairs – knowing and not knowing – but makes it difficult to capture unawareness (Modica and Rustichini, 1994; Dekel et al., 1998). We introduce a question-answer knowledge structure that allows us easily to draw an important distinction between *three* different states: knowing (represented by a question and a particular answer); not knowing, but knowing that one doesn't know (represented by a question and a set of possible answers); and not knowing and not knowing what one doesn't know (represented by the absence of an activated question). This third state corresponds to pure unawareness (Li, 2008), in the sense that an individual is unaware of the question itself and does not even distinguish different possible answers. (In contrast, our question-answer structure does not capture partial unawareness, in the sense of an individual being aware of a question and proper subset of possible answers, but unaware of some other remaining possible answers.) The question-answer structure is consistent with, and could be cast in terms of, a generalized state-space model (e.g., Modica and Rustichini, 1999; Heifetz et al., 2006), but we find the question-answer structure more convenient to use.

The question-answer knowledge structure is intended to reflect human information-processing capabilities. Our cognitive maps of the world are not sets of possible states, each described in exquisite detail to account for all possible consequences of all possible decisions. Instead, people attend to a few relevant aspects of a situation and use limited information to make a broad judgment that can be refined later, if necessary. People tend to set goals and monitor their progress toward them in order to navigate a complex world (Miller et al., 1960; Locke and Latham, 1990; Loewenstein, 1999). In this paper, we advance the idea that the acquisition of knowledge is also goal-oriented. We don't simply seek out information to maximize the data available to us or even to optimize future decisions, but instead tend to seek answers to questions that are either posed to us or that we pose to ourselves. Questions are, therefore, very much like informational goals or reference points. Indeed, focusing on a question that one cannot answer – e.g., a puzzle one cannot figure out – can torment a person and at the same time motivate the search for an answer, much as a high reference point can simultaneously detract from utility and motivate one to strive to reach it.

Beyond the instances of information-seeking and avoidance depicted in Table 1, our framework can acknowledge perhaps the most universal informational phenomenon to which the conventional economic analysis fails to do justice: the search for knowledge and insight. Aristotle in 350 B.C. asserted, "All men by nature desire to know." John Stuart Mill agreed, in his classic *Utilitarianism*, defending the utilitarian approach from critics of his time who argued that the hedonic notion of maximizing pleasure and minimizing pain was dehumanizing. Mill argued that, "it would be absurd that while, in estimating all other things,

[2]Several other papers have also shared the notion that people care about their beliefs (Akerlof and Dickens, 1982; Abelson, 1986; Loewenstein, 1987; Geanakoplos et al., 1989; Asch et al., 1990; Yariv, 2001; Kadane et al., 2008).

quality is considered as well as quantity, the estimation of pleasures should be supposed to depend on quantity alone." Mill then continues with what may be the most famous passage in all of his work: "It is better to be a human being dissatisfied than a pig satisfied; better to be Socrates dissatisfied than a fool satisfied." We too assert that knowledge can be a very real source of utility. A perspective that information derives value solely from its ability to yield material consumption fails to appreciate the most profound benefits provided by information, the knowledge and wisdom it confers.

2 Theoretical Framework

2.1 Cognitive States

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 one's cognitive state, encompassing the attention paid to each of the issues or questions that one is aware of as well as subjective judgments about the possible answers to these questions. While people have preferences about their beliefs (and the attention paid to them), we do not treat beliefs (or attention) as choice variables (as Brunnermeier and Parker (2005) do). People can choose whether or not to acquire information that will influence beliefs, but we assume that one's beliefs, given one's information, are constrained by Bayesian inference.

While there surely is an infinite set of possible states of the world, we assume, realistically we believe, that a person can only conceive of a finite number of questions at any one time. We represent awareness with an array of '*activated*' questions and a remaining set of '*latent*' questions. Activated questions are those that the individual is aware of. Latent questions are those that the individual could become, but is not currently, aware of. The finite subset of questions a person is aware of (i.e., paying at least some attention to) is denoted $\mathcal{Q}$. We label these activated questions as $Q_1, \ldots, Q_m$. A vector of attention weights $\mathbf{w} = (w_1, \ldots, w_m) \in \mathbb{R}^m_+$ indicates how much attention each activated question gets.[3] These attention weights depend on three factors that we designate "*importance*," "*salience*," and "*surprise*." We return to define and discuss these concepts in Section 3.

A question Q_i has a countable set[4] of possible (mutually exclusive) answers $\mathcal{A}_i = \{A_i^1, A_i^2, \ldots\}$.[5] A person may not know the correct answer to a given question, but reasonably 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 2.

Anticipated material outcomes, or prizes, can also be incorporated into this framework. We let X denote a countable set of prizes – i.e., 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

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

[4]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.

[5]We assume that there is no such thing as an answer that is disconnected from a question.

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

Table 2: The question-answer knowledge structure.

as well as about the likelihood of different material outcomes. Denote the space of answer sets together with prizes as $\alpha = \mathcal{A}_1 \times \mathcal{A}_2 \times \cdots \times \mathcal{A}_m \times X$. Thus, given a state of awareness defined by the set of activated questions $\mathcal{Q}$,[6] we represent a person's cognitive state C with a subjective probability measure π defined over α (i.e., over possible answers to activated questions as well as eventual prizes) and a vector of attention weights $\mathbf{w}$. 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[7]). 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.[8]

The formal representation of a cognitive state is depicted in Table 3. Consider, for example, a college

Activated Questions	Possible Answers	Subjective Probabilities*	Attention Weights
Q_1	$\mathcal{A}_1 = \{A_1^1, A_1^2, \ldots\}$	$[\pi_1(A_1^1), \pi_1(A_1^2), \ldots]$	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

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

Table 3: Representation of a cognitive state.

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?" $(0, 1, 2, \ldots)$; "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. The material outcome might be next year's salary, which would also depend on (but not be completely determined by) whether or not she gets tenure. Looking at the ratings will definitively answer the first question and may resolve some, but not all, of the uncertainty surrounding the other issues.

[6]In most cases, we will assume that activation of questions is determined exogenously – i.e., by the environment. We don't model growing awareness (see Karni and Vierø, 2013). We recognize that in some cases an individual's actions may change the environment, in turn leading to changes in activation. Moreover, in some unusual cases, such as the choice to read a whodunit, the individual may even anticipate that an action will cause a question to be activated, albeit without knowing what that question will be. We discuss situations like these in Section 6. Until then, assume a fixed $\mathcal{Q}$.

[7]We discuss entropy in Section 3. 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.

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

2.2 Actions

A decision maker has the possibility of taking actions with two kinds of effects: *informational* actions contribute to subjective judgments about the world by answering a question; and *instrumental* actions affect the chances of receiving various prizes (outcomes). 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. Informational actions affect the subjective probability measure through the conditioning of beliefs on the discovered answer. Instrumental actions affect beliefs directly by changing the distribution over prizes conditional on subjective judgments. Both instrumental and informational actions also impact attention weights through their respective effects on importance and surprise. 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.

At any point in time an individual can be characterized by a prior cognitive state consisting of subjective probability measure π^0 and attention weight vector $\mathbf{w}^0$. Actions, in general, are operators on cognitive states that map to new cognitive states or to distributions over cognitive states. A purely instrumental action acting on the prior cognitive state determines a particular new cognitive state. Typically, it preserves the prior subjective 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 subjective probability measure over $\Delta(\alpha)$ changes from π^0 to $\pi^{A_i} = \pi^0(\cdot|A_i)$.[9] We assume Bayesian updating here, which means that ex ante, before one 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, $\sum_{A_i \in \mathcal{A}_i} \pi^0_i(A_i)\, \pi^{A_i} = \pi^0$. An informational action would decrease expected entropy because conditioning reduces entropy (see, e.g., Cover and Thomas, 1991, pg. 27). New information generates surprise (as formalized in the next section), which changes the attention weights too. Given the prior attention weight vector $\mathbf{w}^0$ based on salience and importance, we let $\mathbf{w}^{A_i}$ denote the new attention weight vector immediately after learning A_i, resulting from surprise at this discovery.

2.3 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. Given that we may think of an informational action as creating a lottery over cognitive states, we make the natural assumptions leading to an expected utility representation in this new domain.[10]

[9] We thus denote a belief with complete certainty in $\mathbf{A} \times x$ as $\pi^{\mathbf{A} \times x}$.

[10] While nonlinear probability weighting or reference-dependent valuation are as plausible for information preferences as for those involving only outcomes, we adopt an expected utility representation as a simplification to make the model tractable and to focus attention on the novel aspects of our framework. Additionally, although it is not our focus in this paper, behavior that has been ascribed to these effects might instead be accounted for by considering cognitive states, rather than material outcomes, as the objects of valuation in an expected utility model. For example, low-stakes risk aversion (Rabin, 2000), typically attributed to

Independence Across Cognitive States

We assume that *there is a complete and transitive preference relation* $\succeq$ *on* $\Delta(\mathcal{C})$ *that is continuous (with respect to an appropriate topology)*[11] *and that satisfies independence*, so there exists a continuous expected utility representation u of $\succeq$ (von Neumann and Morgenstern, 1944).

The assumption here is 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.

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 – e.g., 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.

2.4 Choosing Between Sequences of Actions

The discovery of information following an initial action can change the availability or desirability of subsequent actions. For example, the information in a college professor's teaching ratings could help her decide whether to enroll in a teacher improvement class. A sequence of actions can be analyzed with the convention that an action operator passes through a distribution over cognitive states.[12] Thus, we represent a sequence of actions s acting on a cognitive state $(\pi, \mathbf{w})$ as $s\cdot(\pi, \mathbf{w}) \in \Delta(\mathcal{C})$.

Choice from among a set of sequences of actions $\mathcal{S}$, 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 decision maker in the cognitive state $(\pi, \mathbf{w})$ if $s^* = \arg\max_{s\in\mathcal{S}} u\left(s\cdot(\pi, \mathbf{w})\right)$. We find it useful to define a utility function over cognitive states, contingent on the set of sequences of actions that may subsequently be chosen:

$$U(\pi, \mathbf{w} \,|\, \mathcal{S}) = \max_{s\in\mathcal{S}} u\left(s\cdot(\pi, \mathbf{w})\right). \tag{1}$$

In the example of the professor's teaching ratings, 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 having utility that is conditional on making the optimal choice of one of these subsequent actions.

---

loss aversion, could also be attributed to the discomfort of thinking about uncertainties, and the endowment effect, also commonly thought of as a consequence of loss aversion, could arise if people ask themselves if they are getting a good deal (and care about the answer) (Weaver and Frederick, 2012). Perhaps, if violations of the independence axiom in choice under risk arise only due to utility derived from beliefs, then by incorporating beliefs directly into the utility function we may salvage this seemingly intuitive axiom. Indeed, the fundamental concept of an information gap in our question-answer framework already has elements that are reminiscent of reference-dependence even without building such dependence explicitly into the model.

[11] 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).

[12] 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.

We define the desirability of a sequence of actions s in cognitive state $(\pi, \mathbf{w})$ as

$$D(s \mid \pi, \mathbf{w}) = u\left(s \cdot (\pi, \mathbf{w})\right) - u\left(\pi, \mathbf{w}\right).^{13}$$

The desirability of taking a course of action is therefore the marginal utility relative to the trivial 'action' of doing nothing (not changing the cognitive state).

2.5 The Desire for (or to Avoid) Information

Having assumed independence across cognitive states, it follows directly that the utility of receiving information can be captured as the difference between the expected utility after receiving the information and the ex ante utility before receiving the information. Given a prior cognitive state $(\pi^0, \mathbf{w}^0)$ and a set $\mathcal{S}$ of subsequent sequences of actions available to the decision maker, we define the desire for information answering question Q_i as

$$D_i \;=\; \sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i)\, U\big(\pi^{A_i}, \mathbf{w}^{A_i} \mid \mathcal{S}\big) \;-\; U\big(\pi^0, \mathbf{w}^0 \mid \mathcal{S}\big)\,. \tag{2}$$

Naturally, when this quantity is positive (or negative), an individual seeks (or avoids, respectively) the answer to question Q_i.

Learning the answer to a question has three consequences: 1) the information may affect the value of subsequent actions that may be chosen from $\mathcal{S}$; 2) the information may change the probabilities associated with different answers (the transition from π^0 to π^{A_i}); and the information may change the attention weights (the transition from $\mathbf{w}^0$ to $\mathbf{w}^{A_i}$). We can now identify in Equation 2 three corresponding sources for the desire to acquire or to avoid information: 1) the instrumental value of that information; 2) curiosity; and 3) motivated attention.

Instrumental value, derived from the usefulness of information for affecting future actions (e.g., informing future decisions), has been recognized in economics for a long time. The instrumental value of information answering question Q_i, when the set $\mathcal{S}$ of subsequent sequences of actions is available, is

$$D_i^{\text{IV}} = \sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i)\; \max_{s \in \mathcal{S}} D(s \mid \pi^{A_i}, \mathbf{w}^{A_i}) \;-\; \max_{s \in \mathcal{S}} D(s \mid \pi^0, \mathbf{w}^0). \tag{3}$$

This is the difference between the expected utility gain from subsequent actions after having acquired the information and the utility gain that could be derived from subsequent actions without having this information (see, e.g., Hirshleifer and Riley, 1979).

Curiosity has been recognized by psychologists as the desire for knowledge for its own sake, apart from any benefits that knowledge may confer. We isolate curiosity for the answer to question Q_i as

$$D_i^{\text{C}} = \sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i)\, u\left(\pi^{A_i}, \mathbf{w}^0\right) \;-\; u\left(\pi^0, \mathbf{w}^0\right). \tag{4}$$

This is the gain in utility from updating beliefs, holding attention weights fixed.

[13] The degenerate distributions in $\Delta\left(\mathcal{C}\right)$ correspond to individual states of knowledge. With the standard abuse of notation, we refer to the utility of the degenerate distribution on $(\pi, \mathbf{w}) \in \mathcal{C}$ as $u(\pi, \mathbf{w})$.

Motivated attention to (or avoidance of) information arises from the impact of an informational action on attention. We express this as

$$D_i^{\mathrm{MA}} = \sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i)\ \left(u\left(\pi^{A_i}, \mathbf{w}^{A_i}\right) - u\left(\pi^{A_i}, \mathbf{w}^{0}\right) \right). \tag{5}$$

Putting together these three motives yields the desirability of information answering question Q_i,

$$D_i = D_i^{\mathrm{IV}} + D_i^{\mathrm{C}} + D_i^{\mathrm{MA}}.$$

We address each of these motives, along with the phenomena they produce, in Section 4.

3 Psychological Insights

In this section we introduce a number of specific psychological insights that lead us to specify a utility function that generates a wide range of testable predictions concerning informational phenomena. These insights help us characterize the factors that influence the level of attention paid to a question as well as to identify distinctly the valence of beliefs and the desire for clarity. In Section 4 we detail the specific predictions that follow from the integration of these insights into the formal framework outlined in Section 2.

3.1 Attention

Neuroeconomic research indicates that attention shapes preference (Fehr and Rangel, 2011), and behavioral economists are now recognizing the role attention plays in a variety of economic behaviors (Gabaix and Laibson, 2006; DellaVigna and Pollet, 2009; Bordalo et al., 2012a; Bordalo et al., 2012b; Köszegi and Szeidl, 2013; Bhatia and Golman, 2013). 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 may think of beliefs as having intrinsic value, which is then amplified by these attention weights. Our model (assuming monotonicity with respect to attention weights, as described in the appendix) provides a natural distinction between beliefs that have positive or negative intrinsic value: beliefs are positive specifically when more attention enhances utility and are negative in the opposite case. That is, a person likes thinking about (i.e., putting more attention weight on) *positive beliefs* and does not like thinking about *negative beliefs*.

Here we formalize the concepts of *importance*, *salience*, and *surprise*, all of which, we assume, contribute to attention weight. The importance γ_i of a question Q_i reflects the degree to which one's utility depends on the answer. Thus, for example, for an egocentric, but insecure, individual, the question, "Do other people like me?" is likely to be of great importance because the answer matters to the individual. Salience, distinctly, reflects the degree to which a particular context highlights the question. If, for example, an individual hears that another person was talking about her (with no further details), the question of whether the comments were favorable or not will become highly salient. 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. If, having believed that she was generally well-liked, our 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, as we presently assume, would increase her attention to the question). We denote the surprise associated with a revised belief about question Q_i as δ_i. We assume that *the attention w_i on an activated question Q_i is a strictly increasing function of this question's importance γ_i, its salience σ_i, and the surprise δ_i associated with it*.

Importance

The importance of a question depends on the spread of the utilities associated with the different answers to that question. The degree to which an individual's utility varies with the answers to a question depends both on the magnitude of the utility function and on the perceived likelihood of different answers. Continuing with the example of the question of how well-liked an individual is, one could distinguish two relevant traits: egocentrism – the degree to which the individual *cares* about being well-liked; and insecurity – the dispersion of the individual's subjective probability distribution across possible answers. By our definition of the concept, importance should be positively related to both properties.

Given a particular prior subjective 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

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

where U is the utility function defined in Equation (1). Without specifying the precise form of this function ϕ, we assume only that *it (i.e., importance) increases with mean-preserving spreads*[14] *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, a question is important to the extent that one's utility depends on the answer.[15] Raising the stakes increases importance. On the other hand, if an answer is known with certainty, then by this definition nothing is at stake, so the underlying question is no longer important. While acquiring information will affect the importance of the questions being addressed, it takes time to adapt to news, so there should be some delay. We assume that *the importance of a question is updated only when the new information is incorporated into a new default subjective probability measure*.

Our definition of importance is, admittedly, circular. Importance depends on utility, which in turn depends on the attention weight, but importance also contributes to attention weight. There is, likely, some 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

[14]The concept of a mean-preserving spread is also built into the notion of second-order stochastic dominance.

[15]This definition encompasses many sources of importance. Questions may be intrinsically important, meaning that utility is directly dependent on the answer. Similarly, questions may have implicit importance if one cares about the answer to a correlated question, i.e., if the answer reveals a clue about something else with underlying intrinsic importance. Questions may also be materially important, meaning that the prize correlates with the answer (and utility is dependent on the prize). To make the comparisons concrete, the outcome of a competition between a home team and a divisional rival would be intrinsically important, and the outcome of a preseason tuneup game might be implicitly important for what it reveals about the home team's prospects for the coming year, whereas the outcome of a game on which one has wagered, but otherwise does not care about, would be materially important. In another instance, material importance may be derived from a subsequent decision. For example, the outcome of the preseason game might be materially important if one is deciding whether to bet on the teams' upcoming games.

and utility) will be a fixed point of this composition of functions. 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.

Often a question may be salient despite being unimportant. Continuing the prior example, even if an individual deems others' perceptions of her as unimportant, the question of her popularity might nonetheless be highly salient if the individual was asked, "Do you know what x thinks of you?" Conversely, there are myriad questions that are important by the definition just provided, but which lack salience. There might be numerous people whose opinion of us we would care about and be unsure of, but unless something raises the issue in our mind, we are unlikely to focus on it. 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). Thus, we assume that a question Q_i is activated (i.e., has strictly positive attention weight $w_i > 0$) if and only if it has positive salience $\sigma_i > 0$. Further, 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)*. That is, an increase in importance produces a greater increase in attention for a more salient question.

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. A natural measure of surprise was proposed in a theoretical paper by Baldi (2002) and, in an empirical follow-up investigation (Itti and Baldi, 2009), shown to predict the level of attention paid to information. Incorporating the insights from this line of research, we assume that *when the answer to a particular question Q_j is learned, thereby contributing information about the answers to associated questions and causing their subjective 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*:

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

Surprise is positive with any new information, and is greatest when one learns the most unexpected answer with certainty. However, the feeling of surprise is not permanent. We assume that *when the decision maker adapts and gets used to this new knowledge (formally, when the default subjective probability measure is reset), it is 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 immediately spikes, but in the long run 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 over time this attention weight falls to a lower level. Specifically, when the decision maker adapts and the certain belief is incorporated into the default subjective probability measure, 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.

3.2 Valence and Clarity

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. 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. Clarity refers to preferences between degrees of certainty, independent of the answers one is certain of. We assume that, *ceteris paribus, people prefer to have greater clarity (i.e., less uncertainty or more definitive subjective beliefs).* The aversion that people feel towards uncertainty is reflected in neural responses in the anterior cingulate cortex, the insula and the amygdala (Hirsh and Inzlicht, 2008; Sarinopoulos et al., 2010). 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 central to the phenomenon of curiosity, which we discuss in greater depth in Section 4.2. Similarly, the desire to (partially) control the attention weight on beliefs with positive or negative valence is central to the phenomenon of motivated attention, which we discuss further in Section 4.3.

When valence and clarity pull in opposite directions, it may be the case that people prefer a certain answer to a subjective belief that dominates it on valence or that people prefer uncertainty when it leaves space for better answers. While the preference for clarity violates Savage's (1954) sure-thing principle, we do assume a weaker version of it:

One-Sided Sure-Thing Principle

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 $\mathbf{A}' \neq \mathbf{A}''$.

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

A Measure of Uncertainty

The assumption of a preference for clarity means that there is a preference for less uncertain subjective beliefs. A useful measure of the uncertainty about a particular question is the entropy of the subjective probability distribution over answers (Shannon 1948). The entropy of a subjective (marginal) probability π_i is $H(\pi_i) = -\sum_{A_i \in \mathcal{A}_i} \pi_i(A_i)\, \log \pi_i(A_i)$.[16] At one extreme with maximal entropy, there are many equally likely possible answers; at the other extreme with minimal entropy 0, there is a single answer known for sure.

3.3 A Specific Utility Function

To make precise predictions about preferences for (or to avoid) information, we consider a specific utility function incorporating the preference for clarity and the role of attention weights:

$$u(\pi, \mathbf{w}) = \sum_{x \in X} \pi_X(x) v_X(x) + \sum_{i=1}^{m} w_i \left(\sum_{A_i \in \mathcal{A}_i} \pi_i(A_i) v_i(A_i) \; - H(\pi_i) \right). \tag{6}$$

We represent the value of prize x as $v_X(x)$ and the valence of answer A_i as $v_i(A_i)$. We describe properties (some quite strong and almost certainly not always satisfied) that characterize (and necessarily imply) this utility function in the appendix (Theorem 1). Our primary focus now is exploring the patterns of behavior that this utility function implies.

4 Information Acquisition and Avoidance

As in the standard economic account, the possibility of taking subsequent actions after acquiring information gives useful information instrumental value. Additionally, by incorporating the preference for clarity in the utility function in Equation (6), we are able to capture curiosity. The fact, mentioned in passing in Section 2, that acquiring information decreases expected entropy implies that curiosity (as represented by Equation (4)) is always positive. And finally, by incorporating attention weights in the utility function, we can capture motivated attention to (or avoidance of) information. Whereas curiosity can only motivate an individual to acquire information, motivated attention (as represented by Equation (5)) may at times drive a person's desire to avoid information. We address each of the three motives in turn, before concluding this section with predictions about information acquisition and avoidance that result from integrating these distinct motives.

4.1 The Instrumental Value of Information

The traditional economics of information is solely concerned with the usefulness of information for making future decision, as illustrated by a professor's desire to examine her teaching ratings before deciding whether to enroll in a teacher improvement class. Instrumental value from information that allows one to make a better choice among subsequent actions shows up in Equation (3). In addition, by recognizing the utility of beliefs, Equation (3) incorporates a new source of instrumental value: information can make an intended subsequent action more (or less) attractive. An art lover, for example, might read about an artist before observing his new exhibit at the museum. Discovering an initial piece of (surprising) information about the

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

artist may increase attention on relevant questions about his style, stimulating curiosity about the exhibition and, in turn, leading to a greater utility gain when this curiosity is eventually resolved. Similarly, a person reading a novel might ask a friend not to give away the ending, temporarily avoiding information until it will have the most impact (and thus not ruining a good surprise). By the same token, instrumental actions can also alter the value of information that one intends to acquire. For example, giving away money may have additional value (beyond the direct selfish cost and altruistic benefit) if a person anticipates finding out how the recipient views or uses the gift (Ellingsen and Johannesson, 2008).

4.2 Curiosity

While curiosity has gotten short shrift in the economics literature, there is ample evidence that it is distinct from instrumental value as a motive for acquiring information. There are countless things people want to know despite having no use for. There is a natural inclination to resolve information gaps, even for questions of no importance and even when all possible answers have neutral valence. Of course, curiosity can also complement instrumental value, and people are often more curious about information that promises to be useful.

In a review of the psychological literature on curiosity, one of the authors of this paper (Loewenstein, 1994) introduced the concept of an information gap and proposed that curiosity occurs when an individual becomes aware of a gap in his or her knowledge that could potentially be filled by information. That is, a gap opens when a question becomes *activated* but the answer is not known with certainty. The association of curiosity with an information gap that is attracting attention suggests a natural explanation for the fact that, as Kang et al. (2009) reported, subjects are better able to recall the answers to questions that they have previously reported being curious about. To wit, curiosity results, in part, from increased attention on a question, which should aid memory for the answer. Indeed, Kang et al. (2009) link curiosity to pupil dilation, which is a well-known, reliable measure of attention (Kahneman, 1973). The information gap account suggests that an individual's degree of curiosity will depend on three factors.

The first factor is somewhat obvious: curiosity tends to be stronger about questions that are more *important*. For example, people tend to care a lot about what kind of person they are, and about how others perceive them – even when knowledge of this information has no practical use. So, people will naturally be more curious about information that helps to answer these questions than about information that could help to answer similar questions, but about another person. People also tend to care about material outcomes beyond their control.[17] Returning to the example of a college professor's teaching ratings, suppose the professor is uncertain whether she is a good teacher or whether she will get tenure, so that these questions are important and get a lot of attention. (Being a good teacher may not necessarily correlate with getting tenure, but it nevertheless may be important if she intrinsically cares about it.) If her teaching ratings correlate sufficiently with these important questions, these ratings may be very important as well. In this case, the professor would pay a lot of attention to her ratings and be very curious to see them. On the other hand, if the professor already has tenure and great confidence in her teaching ability, then another semester's ratings may not be so important, and curiosity would be slight.

[17] Recent research showing that people are willing to pay for non-instrumental information relating to anticipated payoffs (Eliaz and Schotter, 2010; Powdthavee and Riyanto, 2014) is consistent with this insight.

Second, and more interestingly, curiosity tends to be an increasing function of the *salience* of the information gap. For example, if a college professor hears her colleagues discussing their teaching ratings, the issue would be more salient to her, and she would be more curious about her own. For another example, consider a student's desire to know the score they obtained on a test. One can imagine a progression of contextual factors that would lead to increasing salience (Table 4) and, thus, increasing attention weight. We

Increasing Salience	not yet taken
↓	taken, but not yet scored
	taken and scored
	taken and scored, and displayed by someone who knows the score

Table 4: The salience of a test.

would thus predict that the student would become successively more curious about the test score after it was taken, after it was scored, and then when the teacher displays the graded tests to the class. As we discuss in Section 4.3, however, curiosity about the test score could be offset by a desire to avoid thinking about the test if one expects a bad score.

The third factor contributing to curiosity is *epiphany* – people are especially motivated to acquire information that has the potential to fill multiple information gaps at once. Most items of information that address one question are also likely to shed light on other questions. For example, the answer to the question "What were my teaching ratings last semester?" might also shed light on the question "Will I get tenure?" as well as "Am I a smart, clever and articulate individual?" The eureka moment of sudden comprehension involves discovering a new piece of information, perhaps insignificant by itself, but which resolves many other questions coherently. If a person became aware of a new question whose answer correlates with many questions she is already asking, she'd have particularly strong curiosity because answering this question would help her answer many other activated questions at the same time.

We can see, by examining Equation (4), how each of these factors – the importance of an information gap, its salience, and the potential for epiphany – drive curiosity in our model. Adopting the utility representation in Equation (6), we see that curiosity (as formalized in Equation (4)) comes from the expected reduction in entropy of uncertain beliefs, weighted by the attention placed on those beliefs.[18] (The terms involving valence cancel out.) Answering a question with greater relevance to other questions (i.e., greater potential for epiphany) offers more potential reductions in expected entropy, so such questions generate stronger curiosity. Similarly, questions that are more important or more salient attract more attention weight, which magnifies the expected entropy reduction, and thus such questions are also subject to greater curiosity. Both arguments rely on the fact that acquiring information necessarily reduces expected entropy (assuming Bayesian updating).

One surprising feature of curiosity discussed in Loewenstein's review is that the pleasure one derives from obtaining information one is curious about often seems incommensurate (on the negative side) with the intensity of the drive to obtain the information. A juicy nugget of gossip is eagerly received and soon

[18]Entropy times attention weight, as built into Equation (6), satisfies Berlyne's (1957b) criteria for a measure of the internal conflict or dissonance in one's cognitive state. Conflict, the potential for surprise, and uncertainty are all drivers of curiosity, and are all related through the concept of information entropy (Berlyne, 1954; Berlyne, 1957a).

forgotten. This property is naturally accommodated by the assumptions that surprise fades over time and that answered questions (or filled information gaps) lose their importance – the "belief resolution effect" discussed in Section 3. The attention weight associated with a particular question initially rises when the definitive answer is learned, but ultimately falls below its prior level. The satisfaction of curiosity will be disappointing to the extent that this drop in attention weight occurs rapidly (as seems likely to be the case) and unexpectedly.

4.3 Motivated Attention

Motivated attention allows desire for information to depend on the valences of possible answers even though one cannot know a priori which answer the information will reveal. In our model, simply getting a definitive answer attracts attention through surprise. If an answer has positive valence, thinking more about it when it is revealed increases utility (even if an answer that good was expected). Conversely, thinking more about an answer with negative valence decreases utility (even if an answer that bad was expected).

Return once more to the example of the college professor's teaching ratings. If the professor is uncertain whether she is a good or a great teacher (i.e., all possible answers have positive valence), then she would enjoy looking at her ratings. On the other hand, if she is uncertain whether she is a bad or an awful teacher (i.e., all possible answers have negative valence), then looking at (and thinking about) her ratings may be unpleasant. Whether she does want to look at her ratings could well depend on whether or not her students applauded on the last day of class. (If the students applauded, her beliefs about her teaching ratings and ability should shift and have higher expected valence, yet with sufficiently many possible answers the shape of her subjective probability distribution may not change much.) Generally, in the short-term it is better to receive news when one suspects that the news is good and worse to receive news when one suspects that the news is bad. The successful teacher is likely to tear open the course evaluations, whereas the unsuccessful teacher is more likely to dispose of the unopened envelope or hide it in a place where it will hopefully be forgotten.

It is straightforward to see in Equation (5) that motivated attention is increasing in the valences of possible answers because the updated attention weights (immediately upon acquiring new information) increase due to surprise. The extra attention weight amplifies the value of newly acquired beliefs, leading to a gain or loss in utility. Naturally, people prefer to think about positive rather than negative situations, so they tend to desire information about questions with high-valence answers and to avoid information about questions with low-valence answers (assuming, as this distinction requires, that utility is separable across certain subsets of questions). For example, most people enjoy opening a gift (an informational action, as distinct from the instrumental action of accepting the gift). On the other hand, most people do not enjoy going to see the doctor for a diagnosis.

The "belief resolution effect" (discussed in Section 3) implies, however, that this ranking of situations is likely to reverse in the longer-term when attention weights on updated beliefs drop because over time people are better able to adapt to definitive than to uncertain conditions. According to our theory, surprise fades and certainty allows one to pay less attention to the bad news because it eventually seems less important. One study found that people with temporary (i.e., potentially reversible) colostomies (a medical procedure in which one's bowels empty into a bag) reported greater levels of happiness (lower levels of misery) right after the procedure than those with permanent ones, but over a short span of time the happiness levels of the two groups reversed; those with permanent colostomies then reported relatively greater levels of happiness (Smith et al., 2009). So, it might be better initially to have definitive good news, and worse to have definitive bad news, but over time the situation is likely to change because people adapt to both good and bad news, when it is definitive. While ignorance may be bliss, a persistent nagging doubt about the possibility of a negative state of affairs, such as a concern that one's child might be taking drugs, tends to be quite unpleasant. Despite the long-term consequences, we expect that people typically avoid confronting issues they don't like thinking about, but we also recognize that people with greater foresight may choose to obtain information about such issues and may ultimately feel better for having done so.

The same situation, but in reverse, occurs for positive information. Research in psychology suggests that uncertainty, e.g., about whether something is true, or about why it might be true, can prolong the pleasure of good news (Wilson et al., 2005). In 3 experiments, Wilson and coauthors induced experimental subjects to experience a positive event (e.g., receive an unexpected gift of a dollar coin) under conditions of certainty or uncertainty (e.g., it was easy or difficult to make sense of the text on the card). Subjects' positive moods lasted longer in the uncertain conditions, though they also found that people seemed to be unaware that this was the case. This lack of awareness suggests, first, that people are most likely to make decisions based on initial reactions (seeking news that clarifies positive beliefs and avoiding news relating to negative beliefs), and, second, that these decisions are unlikely to maximize long-term experienced utility.

To the extent that people *are* aware of the dynamic consequences of certainty and uncertainty for negative and positive outcomes, we should predict that people who are more far-sighted – who discount the future less – will be more prone to resolve uncertainty about negative events so as to 'take the hit' then get on with their lives. That, in fact, has been found – people with low time discounting (as measured by self-reported financial planning horizons) are more likely to undergo cancer screening (Picone et al., 2004). By the same token, we might also predict that people who are more short-sighted will be more prone to resolve uncertainty about positive events, enjoying the momentary pleasure, but shortening its duration.

4.4 Combined Effects of Curiosity and Motivated Attention

We can now integrate the insights that curiosity is stronger for questions that have greater importance, salience, or potential for epiphany and that motivated attention depends on the valence of possible answers. These factors contribute to the desire for information.

Proposition 1 *Suppose utility takes the form of Equation (6), and take as an ancillary assumption that the marginal increase in attention due to surprise δ_j is independent of the salience σ_j and the importance γ_j. Suppose, additionally, that there are no subsequent actions available to the decision maker (so that we can disregard the instrumental value of information). For a given question Q_i, each of the following conditions implies that $\hat{D}_i > D_i$ (i.e., the desire for information answering the question will be increased):*

1. *for some pairwise dependent question Q_{j^*} (i.e., some Q_{j^*} with $\pi_{ij^*} \neq \pi_i \cdot \pi_{j^*}$, perhaps Q_i itself), we change the salience from σ_{j^*} to $\hat{\sigma}_{j^*} > \sigma_{j^*}$, while maintaining the importance of all pairwise dependent questions Q_j, $\hat{\gamma}_j \geq \gamma_j$;*

2. *we transform π to $\hat{\pi}$ by changing some prize $x^* \in X$ to $\hat{x}^*$ such that for all pairwise dependent questions Q_j, $\hat{\gamma}_j \geq \gamma_j$ with at least one such inequality strict;*

3. *we transform π to $\hat{\pi}$ by changing some pairwise dependent answer $A^*_\nu \in \mathcal{A}_\nu$ (for which $\pi_{i\nu}(A_i, A^*_\nu) \neq \pi_i(A_i) \cdot \pi_\nu(A^*_\nu)$) to $\hat{A}^*_\nu$ such that $v_\nu(\hat{A}^*_\nu) > v_\nu(A^*_\nu)$ and for all pairwise dependent questions Q_j, $\hat{\gamma}_j \geq \gamma_j$; or*

4. *we change a set of beliefs π to $\hat{\pi}$ such that for some question Q_ν with $v_\nu(\varpi_\nu) \geq 0$ for all $\varpi_\nu \in \Delta(\mathcal{A}_\nu)$, $\hat{\pi}_{i\,\nu} \neq \hat{\pi}_i \cdot \hat{\pi}_\nu$ and $\pi = \hat{\pi}_{-\nu} \cdot \hat{\pi}_\nu$, and for all pairwise dependent questions Q_j, $\hat{\gamma}_j \geq \gamma_j$;*

Condition 1 in Proposition 1 implies that increasing the salience of a question will increase the desire for information addressing it (i.e., answering a related question or perhaps the given question itself), holding all else equal. Condition 2 implies that increasing the (material) importance of a question (by changing the prizes that may be received, depending on the answer) will also increase the desire for information addressing it. Condition 3 implies that changing a relevant answer to one with higher valence, while not decreasing the (intrinsic/implicit) importance of related questions, will do the same. Finally, Condition 4 implies that increasing the number of related questions, about which beliefs are necessarily positive or at least neutral, also has the same effect. The proof is in the appendix.

An immediate implication of this result (Condition 4) is that a single independent question with uniformly non-negative valence answers attracts a positive desire for information, as this desire has necessarily increased from none at all.

Corollary 1 *Suppose, as in Proposition 1, that utility takes the form of Equation (6) and that there are no subsequent actions available to the decision maker. If belief about question Q_i is independent of other beliefs, $\pi = \pi_{-i} \cdot \pi_i$, and only answers with non-negative valence are considered possible, $v_i(A_i) \geq 0$ (and, of course, $\pi_i(A_i) < 1$) for all $A_i \in \text{supp}(\pi_i)$, then information answering this question would be sought, $D_i > 0$.*

Corollary 1 tells us that when answering a question poses no threat to utility, as would be true for a purely 'intellectual' question (e.g., whether a particular tree is an oak or an elm), people generally want the information. On the other hand, when acquiring information might lead to negative beliefs, individuals may choose to avoid this information.

Conditions for Information Avoidance

Empirical studies have revealed strong evidence consistent with the idea that people tend to seek out information likely to confirm suspicions that their objective situation is favorable, and to avoid information most likely confirming that their objective situation is unfavorable.[19] For example, willingness to pay for an

[19]In psychology experiments investigating the "observing response," people typically seek out information that signals upcoming rewards, but subjects sometimes desire information about bad news, while other times they prefer not to receive such information (Lieberman et al., 1997; Fantino and Silberberg, 2010). Similarly, even rats, pigeons, and monkeys persistently observe reward-indicating cues even when their observations have no effect on the reward rate (Prokasy, 1956; Hendry, 1969; Bromberg-Martin and Hikosaka, 2009), while animals generally tend to avoid bad news (see, for example, Jenkins and Boakes, 1973). Of course, animals exhibiting a conditioned response to the pleasures and pains of information are not aware, as some human subjects may be, of the belief resolution effect.

assessment of one's IQ or beauty (relative to others) increases as one's subjective prior belief about this assessment becomes more favorable (Eil and Rao, 2011; Möbius et al., 2011; Burks et al., 2013). Conversely, willingness to pay to avoid testing for a herpes infection is greater for the more dreaded type 2 infection than for type 1 (Ganguly and Tasoff, 2014). That is, as we predict in Proposition 1 (Condition 3), the desire for information increases as the valence of anticipated outcomes increases. Benabou and Tirole's (2002) model of self-confidence and Köszegi's (2006) model of ego utility both make the opposite prediction. They predict that people would have greater desire for information about themselves when they hold negative beliefs about themselves than when they hold positive beliefs about themselves because information may prompt an individual to change a prior belief. While the logic is intuitive, the research just reviewed suggests that this is not typically the case.

Also consistent with our prediction, empirical research on the 'ostrich effect' shows that holders of portfolios who have internet access tend to look up the value of their portfolio – figuratively to "shake their piggy-bank" – when markets are up, but not when they are down (Galai and Sade, 2006; Karlsson et al., 2009; Cai and Meyer, 2013). Karlsson et al. propose an account of the ostrich that relies on the insight that knowing has greater impact on utility than merely suspecting. Similarly, our account follows from the assumption that learning the definitive answer to a question raises the attention weight placed on that question, at least initially. Over the long run, knowing an answer with certainty may allow the attention weight to diminish, but when the answer is first discovered definitively, there is an immediate boost in attention, and this tends to dominate decision making. More generally, our assumption that surprise contributes to attention weight means that whenever the judged subjective probability of answers to a question changes, this updated belief is accompanied by an immediate boost in attention weight, which gives the resulting belief more impact on utility.

Thus, acquiring information to more clearly resolve negative beliefs involves a tradeoff. While information would reduce expected entropy and thereby increase utility through improved clarity, it may also lead to increased attention weight placed on beliefs that decrease utility – i.e, that operate through valence. When beliefs are sufficiently negative, a person may prefer to avoid information. Thus, we predict that as the intrinsic valence of a negative belief gets worse, a typical person (and especially one who is short-sighted) will be less inclined to obtain relevant information. The findings that women with breast cancer symptoms which are getting worse delay longer in seeing a physician than those whose symptoms are steady or disappearing (Caplan, 1995), as do women who had a family member with breast cancer (Meechan et al., 2002), support this prediction.

The pattern of behavior predicted by the model is a clear desire for information relating to positive or neutral beliefs but a desire for or against information with negative repercussions depending on the tradeoff between greater clarity on the one hand and greater attention weight on negative valence on the other. Notably, the tradeoff in deciding whether to acquire information relating to negative beliefs may depend on the prior attention weight. If the marginal increase in attention due to surprise δ_i is independent of the salience σ_i and importance γ_i, then it follows from our utility representation (in Equation (6)) that as the salience or the importance of a question increases, the threshold at which a person prefers to avoid information shifts to increasingly negative beliefs. (See Condition 1 in Proposition 1.) For example, subjects in two laboratory studies were more likely to obtain information about an unchosen gamble (despite the possibility of regret) when this gamble was made more salient, either by providing a clue about the outcome or by determining the outcome but keeping it hidden from the subjects (van Dijk and Zeelenberg, 2007). Similarly, we might expect a person to choose not to obtain a costless medical test so that she may avoid thinking about the possibility that she is sick, but if the test had been conducted and the doctor knew the results, the patient might then prefer to be informed of them. If we weaken the ancillary assumption to allow attention weight to have increasing differences in (δ_i, σ_i) and in (δ_i, γ_i) (as seems more intuitive), then we lose predictive power about the threshold for the ostrich effect, but we can still conclude that for questions associated with positive or at least neutral beliefs, the desire for information should get stronger as the salience or the importance of a question increases.

4.5 Example: Genetic Testing

We can apply our framework to a decision whether to get a medical test, for example, and it gives us a new perspective on Oster et al.'s (2013) findings about the propensity for genetic testing for Huntington's Disease (HD). We begin by describing an individual's cognitive state before getting tested. The activated question that we will focus on is Q, "Do I have the HD gene?" The answer to this question has implications for a wide range of material outcomes, but we might summarize the relevant material outcomes by lifespan T and consumption stream $C(t)$, i.e., we make the gross oversimplification $X = \{T, C(t)\}$. Of course, an individual's lifespan and future consumption are both uncertain (and may even depend in part on future choices, such as when to seek medical attention or when to retire), so the individual has prior belief π about the probability of having the HD gene along with various possible lifespans and consumption streams. (The probability of having the gene is $p = \pi_Q(\text{yes})$ and the probability of not having the gene is $1-p = \pi_Q(\text{no})$, and the probability distribution for lifespan T and consumption stream $C(t)$ is dependent on that answer.) But, while the answer to the activated question Q and the material outcomes X are both uncertain, we have assumed that the individual is aware of the question about the HD gene, whereas she may not be thinking specifically about her lifespan or future consumption. The uncertainty about whether she has the HD gene presents an information gap, and she pays attention to it. Factors that affect the amount of attention weight w on this question (i.e., its salience, its importance, and surprise following new information) then affect whether (and at what cost) genetic testing will be pursued.

Essentially, getting genetic testing changes the individual's cognitive state. With probability p the new cognitive state is defined by belief π^{yes} (in which the probability of having the gene is 1 and the probability distribution over possible lifespans and future consumption streams is updated accordingly) and attention weight w^{yes} (which is an increase in attention commensurate with the surprise associated with the change in belief). With probability $1-p$ the new cognitive state is defined by belief π^{no} and attention weight w^{no}, analogously. We posit that the desirability of the genetic testing is the expected utility of this new cognitive state minus the utility of the prior cognitive state. This change in utility can be decomposed into three parts, which we call instrumental value, motivated attention, and curiosity.

The instrumental value of genetic testing refers to the utility of future choices conditioned on knowing the test results minus the utility of future choices made without knowing the test results. Even though there is no cure for the disease, knowing that one has it has a significant impact on decisions such as whether to have children, when to retire, how much to save, how to invest, and whether to get or stay married (Oster et al., 2013). Standard economic arguments show that additional information can only improve decision making. Thus, the instrumental value of genetic testing is necessarily positive. The fact that genetic testing is rare despite little economic cost (Oster et al., 2013) suggests that one of the other sources of utility (motivated attention, we believe) is negative and significant. However, the instrumental value of medical testing cannot be ignored. If there were a medical treatment that would cure HD, genetic testing would no doubt be commonplace precisely because the information would be so instrumentally valuable in determining whether treatment was necessary.

Motivated attention to avoid genetic testing refers to the expected loss in utility associated with paying more attention to the belief about the HD gene in the cognitive state that arises after finding out the test result than in the prior cognitive state. Any change in belief attracts attention due to surprise. The less likely the individual considers having the HD gene to be, the more surprising it would be if the test does indeed reveal the gene. A "positive" test result would lead to a very negative belief, and having to think more about this negative belief would be very unpleasant. On the other hand, finding out that one does not have the gene would be a relief (perhaps somewhat positive if one's reference point is prior expectations), but the belief might better be characterized as lacking negative valence rather than being intrinsically positive. The ex-ante expectation is that the new cognitive state will have lower utility because it will possibly involve thinking more about the unpleasant state of actually having the HD gene. (The new cognitive state will in the long run (after surprise wears off and importance diminishes with certainty) actually involve less thinking about having (or not having) the HD gene, but people are rarely sophisticated enough to foresee such adaptation. Moreover, the expectation of a loss in utility may not even be a conscious expectation but could arise as a learned response to situations in which one may find out bad news.) This is the key reason, we suggest, that genetic testing for the HD gene is rare.

Even though the overall level of genetic testing is infrequent, the rate of testing increases after symptoms of HD pop up (Oster et al., 2013). Oster et al. interpret this pattern as evidence of a correlation in the ex-ante risk of having HD and the propensity to get the test. Such a correlation could be accommodated in our model (because higher ex-ante probability of having HD implies that a positive test result would be less surprising and would thus lead to a smaller boost in attention on the bad news), but is not necessarily predicted by the model (because higher ex-ante risk also implies that bad news is more likely, so we have a countervailing effect as well). More fundamentally, though, the information-gap framework gives us a new perspective in which attention matters as much as probabilities, and this perspective calls into question whether the pattern of increased testing after symptoms arise really has to do with the ex-ante probability of having HD.

Observing a symptom of HD (or its absence) is itself an instance of acquiring information, which affects both the perceived probability of having HD and the attention to that possibility. In our framework, we could model the daily opportunities to observe symptoms (or the lack thereof) as a series of activated questions, Q_i, "Do I have a symptom on day i?". Having symptoms is obviously highly correlated with having the disease, but symptoms take time to manifest, so on any particular day i the probability of answering "yes" to question Q_i is low. This means that when a symptom pops up, it will produce a significant increase in the probability of having HD and it will be quite surprising, thus attracting extra attention to question Q.

The extra attention could have two consequences: First, if there is a diminishing marginal impact of surprise on attention, it would weaken the marginal impact of additional surprise and thus weaken the impact of motivated attention as a reason to avoid testing (i.e., a person might think, "now that I'm worried I might have the disease, I can't avoid thinking about it, so I might as well find out"). Second, it would increase curiosity to find out if one has the disease (i.e., people would find it uncomfortable to wonder whether they have the disease, and the more they have to think about not knowing, the more curious they would be to find out). After observing a series of days without symptoms, by the same logic, the probability of having HD will have gradually decreased pretty significantly and none of these observations will be very surprising; so there would not be a large increase in attention to question Q. While a correlation in probability of having HD and propensity to test would suggest that patients in this situation exhibit even lower rates of testing, the data show no systematic variation in testing rates as asymptomatic individuals age (Oster et al., 2013).

The hypothesized belief-resolution effect offers us an additional testable prediction as well. If the relationship between observing symptoms and getting tested is due to changes in attention rather than changes in beliefs about the probability of having the gene, then it follows that individuals forced to wait a period of time after discovering symptoms before they could get tested (i.e., individuals who could adapt to the change in their circumstances) would be less inclined to get tested.

5 Risk and Ambiguity Preference

Section 4 shows 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 aversion are complex topics, and we develop these applications in depth in a companion paper (Golman and Loewenstein, 2014). Here, we provide a broad outline of how the model can be applied in these domains.

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 a second factor: 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, especially in comparison to a team their home team is playing against.

Decision making involving uncertainties that are ambiguous is similar to the case with known risks, but with an additional wrinkle: with ambiguity, there are additional information gaps. 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 paper on the topic.

6 Desire for Wisdom

Section 4.2 discussed a ubiquitous motive to acquire information, curiosity. Our model conceives of curiosity as a manifestation of the desire for clarity, a utility gain that comes from filling an information gap apart from the attentional effects of acquiring this information. The fundamental insight is that knowledge has intrinsic value. Flipped around, we could say that uncertainty has intrinsic costs. The comparison we have made here is between knowing and not knowing. But another comparison is also of interest, albeit harder to investigate empirically – the difference between awareness and unawareness.

While we cannot easily give a person the choice whether or not to become aware of a question, we can at least introspect. For those of us who agree with John Stuart Mill that it would be better to be Socrates than a happy pig, we would posit that awareness of meaningful questions is a source of utility. Equation (6), the utility function which represents preferences between cognitive states given a fixed set of activated questions $\mathcal{Q}$, might be augmented with a term $v_{\mathcal{Q}}(\mathcal{Q})$ capturing the intrinsic value of awareness of particular issues.

Wisdom, the combination of awareness and clarity,[20] is, or at least tends to be, preferable to ignorance. We of course must allow exceptions if we are serious that beliefs have valence that may be negative. The

<table>
<tr><th>Question</th><th>Answer</th><th>Belief</th><th colspan="2"></th></tr>
<tr><td>Latent</td><td>–</td><td>Unawareness</td><td rowspan="3">↓ Awareness
↓ Clarity</td><td rowspan="3">↓ Wisdom</td></tr>
<tr><td rowspan="2">Activated</td><td>Unknown</td><td>Uncertainty</td></tr>
<tr><td>Known</td><td>Certainty</td></tr>
</table>

Table 5: Wisdom, the combination of awareness and clarity.

popular adage that "ignorance is bliss" expresses concern for the negative beliefs that awareness may entail. However, in many natural situations, a person may reasonably anticipate that newfound awareness will bring about neutral or even positive beliefs. In such contexts, information and awareness may be simultaneously acquired. For example, a bird-watcher typically would strictly prefer to learn the name of a previously unnoticed songbird rather than to remain unaware of its existence. Curiosity is behind the desire to catch the name upon becoming aware of the bird's existence, even though the particular name does not really matter, but utility from awareness implies that opening, and then immediately closing, an aversive information gap need not be zero sum. Rather, discovering the new bird's name, acquiring both the question and the definitive

[20] We are aware that this may not be the most common usage of the word *wisdom*, but the distinction between knowledge acquired from a state of uncertainty and knowledge acquired from a state of unawareness is rarely made explicit. The term, "wisdom" seems to adequately capture this distinction if we think of a wise man or woman as not only having the right answers, but also asking the right questions.

answer, produces a net positive utility gain, which is what we designate, in the context of our model, the *utility of wisdom*. We find the desire for wisdom in individuals' varied pursuits of insight and expertise, from a naturalist's passion for identifying flora and fauna to a fan's thirst for new baseball statistics or a connoisseur's discriminating taste for wine.[21]

Conventional economic theory can certainly accommodate choices to devote significant time, money, and effort to developing expertise that is unlikely to confer any material benefits. One could posit, for example, that the consumption experience of the wine connoisseur is different from that of the novice – i.e., that, in effect, they are consuming different wine, even if the label on the bottle is the same. However, while such an approach could, in principle, accommodate almost any observed pattern of preference for wisdom, it seems much more parsimonious to accept awareness and clarity both as direct sources of utility. We would accept that, for example, a wine drinker would prefer to know whether she was drinking a merlot or a shiraz even if she were indifferent between the two wines. She would like to know how those two wines differed in taste, even if it did not help her to make better choices between wines or provide any kind of grist for bragging about her knowledge.

7 Conclusion

In this paper, we propose a framework for making sense of a wide range of phenomena involving the demand for, or in some cases desire to avoid obtaining, information. The standard account of the economics of information, which assumes that information is only desired to the extent that it enhances decision making, leaves out many, if not most, of the diverse reasons why people seek out or avoid information, including pure curiosity and the desire to savor good news and avoid bad news. Recent work by economists has addressed some of these motives (e.g., Caplin and Leahy, 2001; Benabou and Tirole, 2002; Brunnermeier and Parker, 2005; Köszegi, 2006; Benabou, 2013), but the framework proposed here is, to the best of our knowledge, the first to integrate a wide range of these motives in a unified theory.

Although our framework introduces an extensive new apparatus, including the concepts of questions, answers and attention weights, we find it to be quite useful. Within this framework we develop a model that provides a range of testable implications. We summarize these predictions here:

H1 Non-instrumental information tends to be desired whenever it will result in non-negative beliefs, i.e., beliefs an individual does not mind thinking about. Such information tends to be more desirable if it pertains to a greater number of activated questions, to more important questions, to more salient questions, or to questions with uniformly higher-valence answers.

H2 Information that may result in negative beliefs may be avoided, but increasing the salience of a question, increasing its material importance, or uniformly increasing the valence of the answers makes avoidance of information less prevalent.[22]

H3 Individuals who discount the future less should be more likely to expose themselves to information

[21] Lab studies also find that people prefer environments which seem to stimulate new questions and promise to provide relevant information (Kaplan, 1992).

[22] The hypothesis that increasing the salience or the material importance of a question weakens the ostrich effect relies on the ancillary assumption that the effect of surprise on attention is independent of salience and of importance. The ostrich effect could strengthen if salience or importance amplifies surprise.

relating to negative beliefs.[23]

H4 Anticipation that receipt of information will occur, especially in a context that makes it highly salient, motivates people to invest (time, effort, or money) in increasing its expected valence.

The model also has implications about the hedonic consequences of information acquisition. These implications could in principle be tested if we had measures of hedonic states, which could take the form of, for example, self reports, brain measurements, facial coding, or even physiological measurements.

H5 To the degree that people do not anticipate the decline in attention after learning an answer (the belief resolution effect), satisfying curiosity is disappointing; the initial motivation to gain the information is disproportionate to the pleasure gained from it.

H6 Acquiring information relating to negative beliefs actually improves long-term well-being. In the case of positive beliefs, resolving uncertainty may actually shorten the duration of the enjoyment of the belief.

H7 If one can anticipate that a latent, meaningful question has non-negative valence answers, then activating the question and learning the answer leaves one better off than not being aware of the question in the first place.

Our theory provides a perspective that makes sense of a wide range of informational phenomena that have already been documented but not adequately explained, such as the ostrich effect for unpleasant information. In addition, it provides a range of novel testable predictions, such as that people will be willing to pay for non-instrumental information specifically when it addresses salient information gaps – e.g., when they become aware they are missing information that is immediately available or possessed by another person in their immediate proximity. In future research we, and we hope other researchers, will test these predictions empirically, refine the proposed theory, and think up other testable implications.

[23]While we have not built time discounting explicitly into our formal model, it would be straightforward to consider the utility of distinct cognitive states in two distinct periods after discovering information – both before and after adapting to it. We could represent the desire for information as the expected utility gain (or loss) associated with discovering the information plus a discounted expected utility gain (or loss) associated with adapting to the information (as surprise fades and importance is updated). The discount factor here should include both time discounting and an indicator function for awareness that adaptation occurs. In our formal model, we have in effect assumed this discount factor to be zero, as if people generally are unaware of adaptation. Hypothesis H3 recognizes that such an assumption is extreme.

Appendix

Properties

The utility function in Equation (6) satisfies the following seven properties.

Independence Across Prizes

In Section 2 we assumed independence across cognitive states. Independence might extend, as in traditional models, to material outcomes, holding beliefs constant.[24]

P1 Holding the rest of the cognitive state constant, the preference relation satisfies independence across prizes if $u(\pi^{\mathbf{A}}, \mathbf{w}) = \sum_{x \in X} \pi_X^{\mathbf{A}}(x)\, u(\pi^{\mathbf{A} \times x}, \mathbf{w})$.

Property (P1) implies belief-dependent expected utility over lotteries that are independent of beliefs about the world. If we also were to assume belief-independent utility for prizes, then we would gain the ability to reduce compound lotteries consisting of horse races as well as roulette lotteries (Anscombe and Aumann 1963) to single-stage lotteries. However, we believe it is often the case that utility is belief-dependent. We might say that a decision maker often has a horse in the race.

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.

P2 A utility function satisfies additive separability between questions if $u(\pi, \mathbf{w}) = u_X(\pi_X) + \sum_{i=1}^{m} u_i(\pi_i, w_i)$.[25]

Property (P2) 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.

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 $i \neq j$ 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 $\pi \in \Delta(\alpha)$.

[24]The caveat is important because in many realistic situations, material outcomes are associated with beliefs that individuals care about. For example, someone betting on a football game might care about which team wins in addition to caring about whether the bet wins.

[25]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 $\mathbf{w}_{-\tilde{\mathcal{Q}}}$ contains the remaining components of $\mathbf{w}$.

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 marginal 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 $u_i(\pi_i', w_i) - u_i(\pi_i'', w_i) = w_i \left(v_i(\pi_i') - v_i(\pi_i'') \right)$. 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 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, i.e., 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 subjective probabilities of the hypothetical conditions, then we say that the utility function satisfies the reduction of compound questions property. Figure 1 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 $A_i \in \mathcal{A}_{i_h}$,

$$v_i(A_i) = v_j(\tau(\mathcal{A}_{i_h})) + v_{i_h}(A_i) \;\text{ and }\; \pi_i(A_i) = \pi_j(\tau(\mathcal{A}_{i_h})) \cdot \pi_{i_h}(A_i),$$

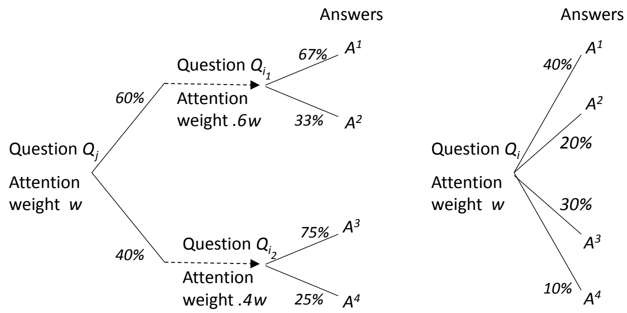

Figure 1: Decomposition of a compound question.

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 subjective 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, i.e. $v_i(A_i) = v_i(A'_i)$ for all A_i and $A'_i \in \mathcal{A}_i$, then for any π 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 $v_i(\pi'_i) > v_i(\pi_i)$.

Utility Representation

Theorem 1 *If the properties P1-P7 are satisfied, then*

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

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 $H(\pi_i) = -\sum_{A_i\in\mathcal{A}_i}\pi_i(A_i)\,\log\pi_i(A_i)$. ∎

Proof of Proposition 1

Conditions 1 and 2 imply that attention weight has been made stronger on some pairwise dependent question Q_{j^*}, $\hat{w}^0_{j^*} > w^0_{j^*}$, and no weaker on other pairwise dependent questions Q_j, $\hat{w}^0_j \geq w^0_j$. Similarly, condition 3 implies that the valence of some answer has increased while the attention weight on all pairwise dependent questions has not decreased. We can consider all three of these cases together, being careful to distinguish (if and) how $\hat{\pi}$ differs from π in each case.

We can apply properties that characterize the utility function. First, using label independence, we define a transformed value with $\hat{v}_X(x^*) = v_X(\hat{x}^*)$ under condition 2 and $\hat{v}_\nu(A^*_\nu) = v_\nu(\hat{A}^*_\nu)$ under condition 3, allowing us to maintain $\hat{\pi} = \pi$. Using Equation (2) we then write

$$\hat{D}_i - D_i \;=\; \sum_{A_i \in \mathcal{A}_i} \pi^0_i(A_i)\; \left(\hat{u}(\pi^{A_i}, \hat{\mathbf{w}}^{A_i}) - u(\pi^{A_i}, \mathbf{w}^{A_i})\right) \;-\; \left(\hat{u}(\pi^0, \hat{\mathbf{w}}^0) - u(\pi^0, \mathbf{w}^0)\right), \tag{7}$$

where the terms representing instrumental value have vanished by assumption. We expand the utility functions according to Equation (6) and group terms in a utility difference

$$\hat{u}(\pi, \hat{\mathbf{w}}) - u(\pi, \mathbf{w}) \;=\; \pi_X(x^*)\left(\hat{v}_X(x^*) - v_X(x^*)\right) + \;\hat{w}_\nu\, \pi_\nu(A^*_\nu)\left(\hat{v}_\nu(A^*_\nu) - v_\nu(A^*_\nu)\right) + \\ \sum_{j=1}^{m}(\hat{w}_j - w_j)\Bigg(\sum_{A_j \in \mathcal{A}_j} \pi_j(A_j) v_j(A_j) \;-H(\pi_j)\Bigg).$$

The ancillary assumption that the marginal increase in attention due to surprise δ_j is independent of the salience σ_j and the importance γ_j tells us that $\hat{w}^{A_i}_j - \hat{w}^0_j = w^{A_i}_j - w^0_j$ or equivalently $\hat{w}^{A_i}_j - w^{A_i}_j = \hat{w}^0_j - w^0_j$. This allows us to extract a common factor of $\hat{w}^0_j - w^0_j$ in the last term of the expansion of Equation (7):

$$\hat{D}_i - D_i \;=\; \left(\sum_{A_i \in \mathcal{A}_i} \pi^0_i(A_i)\, \pi^{A_i}_X(x^*) \;-\; \pi^0_X(x^*)\right)\left(\hat{v}_X(x^*) - v_X(x^*)\right) \;+ \\ \left(\sum_{A_i \in \mathcal{A}_i} \pi^0_i(A_i)\; \hat{w}^{A_i}_\nu\, \pi^{A_i}_\nu(A^*_\nu) \;-\; \hat{w}^0_\nu\, \pi^0_\nu(A^*_\nu)\right)\left(\hat{v}_\nu(A^*_\nu) - v_\nu(A^*_\nu)\right) \;+ \\ \sum_{j=1}^{m}(\hat{w}^0_j - w^0_j)\left(\sum_{A_i \in \mathcal{A}_i} \pi^0_i(A_i)\left(\sum_{A_j \in \mathcal{A}_j}\left(\pi^{A_i}_j(A_j) - \pi^0_j(A_j)\right) v_j(A_j) \;-H(\pi^{A_i}_j) + H(\pi^0_j)\right)\right).$$

We now simplify by applying the law of total probability on each line. The first line vanishes entirely, and the second and third lines reduce to

$$\hat{D}_i - D_i \;=\; \sum_{A_i \in \mathcal{A}_i} \pi^0_i(A_i)\, \pi^{A_i}_\nu(A^*_\nu)\left(\hat{w}^{A_i}_\nu - \hat{w}^0_\nu\right)\left(\hat{v}_\nu(A^*_\nu) - v_\nu(A^*_\nu)\right) \;+ \\ \sum_{j=1}^{m}(\hat{w}^0_j - w^0_j)\left(H(\pi^0_j) - \sum_{A_i \in \mathcal{A}_i} \pi^0_i(A_i) H(\pi^{A_i}_j)\right).$$

In conditions 1 and 2, $\hat{v}_\nu = v_\nu$, so the first line vanishes. Condition 3 specified that $v_\nu(\hat{A}^*_\nu) > v_\nu(A^*_\nu)$, and because surprise can only increase attention weight we know that $\hat{w}^{A_i}_\nu - \hat{w}^0_\nu \geq 0$, with the inequality strict for some A_i (specifically, for the A_i satisfying $\pi_{i\nu}(A_i, A^*_\nu) \neq \pi_i(A_i) \cdot \pi_\nu(A^*_\nu)$). Thus, in condition 3 the sum in the first line is strictly positive. Conditioning on the answer A_i strictly decreases (in expectation) the entropy of the belief about a pairwise dependent question Q_{j^*}, i.e., $H(\pi^0_{j^*}) - \sum_{A_i} \pi^0_i(A_i) H(\pi^{A_i}_{j^*}) > 0$. With $\hat{w}^0_{j^*} > w^0_{j^*}$ in conditions 1 and 2, this second sum is strictly positive. (In condition 3, we know only that it is non-negative because the latter inequality is weak.) Thus, in all three of these conditions, $\hat{D}_i > D_i$.

We now turn to condition 4. It specifies that importance of pairwise dependent questions does not decrease. We have just shown that increased importance of some pairwise dependent question can only increase the desire for information. We now consider the case that importance, and thus the prior attention weight $\mathbf{w}^0$, has not been changed by the transformation $\pi \to \hat{\pi}$ specified by condition 4. Recognizing that $\hat{\pi}^0_j = \pi^0_j$ for all j, we have $u(\hat{\pi}^0, \mathbf{w}^0) = u(\pi^0, \mathbf{w}^0)$ (again using Equation (6)), and the change in the desire for information simplifies as

$$\hat{D}_i - D_i \;=\; \sum_{A_i \in \mathcal{A}_i} \pi^0_i(A_i)\; \left(u(\hat{\pi}^{A_i}, \hat{\mathbf{w}}^{A_i}) - u(\pi^{A_i}, \mathbf{w}^{A_i})\right). \tag{8}$$

Only the updated belief about question Q_ν (conditioning on A_i) and the surprise associated with this belief differ under the transformation in condition 4, so

$$\begin{aligned} u(\hat{\pi}^{A_i}, \hat{\mathbf{w}}^{A_i}) - u(\pi^{A_i}, \mathbf{w}^{A_i}) \;=\; & \left(\hat{w}^{A_i}_\nu - w^{A_i}_\nu\right) v_\nu(\hat{\pi}^{A_i}_\nu) \;+ \\ & w^{A_i}_\nu \left(\sum_{A_\nu \in \mathcal{A}_\nu} \left(\hat{\pi}^{A_i}_\nu(A_\nu) - \pi^{A_i}_\nu(A_\nu)\right) v_\nu(A_\nu) \;\; - H(\hat{\pi}^{A_i}_\nu) + H(\pi^{A_i}_\nu) \right). \end{aligned}$$

Plugging this into Equation (8) and simplifying with the law of total probability, we obtain

$$\hat{D}_i - D_i \;=\; \sum_{A_i \in \mathcal{A}_i} \pi^0_i(A_i)\; \left[\left(\hat{w}^{A_i}_\nu - w^{A_i}_\nu\right) v_\nu(\hat{\pi}^{A_i}_\nu) \;+\; w^{A_i}_\nu \left(H(\pi^{A_i}_\nu) - H(\hat{\pi}^{A_i}_\nu)\right)\right] \tag{9}$$

We know $\hat{w}^{A_i}_\nu \;\geq\; w^{A_i}_\nu$ because there may be surprise about question Q_ν after learning A_i when these questions are pairwise dependent, but there is no surprise about Q_ν when these questions are independent, and surprise only increases attention weight. Moreover, condition 4 specified that $v_\nu(\cdot) \geq 0$, so the first term inside the brackets in Equation (9) is nonnegative. When the questions are pairwise independent, conditioning on A_i does not change the belief about Q_ν, so $\pi^{A_i}_\nu = \pi^0_\nu = \hat{\pi}^0_\nu$. Conditioning on the answer A_i strictly decreases (in expectation) the entropy of the belief about a pairwise dependent question, so

$$\sum_{A_i \in \mathcal{A}_i} \pi^0_i(A_i) H(\hat{\pi}^{A_i}_\nu) \;<\; H(\hat{\pi}^0_\nu) \;=\; \sum_{A_i \in \mathcal{A}_i} \pi^0_i(A_i) H(\pi^{A_i}_\nu).$$

Thus, returning to Equation (9), we conclude $\hat{D}_i > D_i$. ∎

References

Abelson, R. (1986). Beliefs are Like Possessions. *Journal for the Theory of Social Behavior* 16 (3), 223-250.

Akerlof, G., Dickens, W. (1982). The Economic Consequences of Cognitive Dissonance. *American Economic Review* 72 (3), 307-319.

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 (1), 54-64.

Asch, D., Patton, J., Hershey, J. (1990). Knowing for the Sake of Knowing: The Value of Prognostic Information. *Medical Decision Making* 10 (1), 47-57.

Aumann, R. (1976). Agreeing to Disagree. *Annals of Statistics* 4, 1236-1239.

Baldi, P. A Computational Theory of Surprise. (2002). In *Information, Coding, and Mathematics*, M. Blaum, P. Farrell, H. van Tilborg (Eds.) Norwell, MA: Kluwer.

Benabou, R. (2013). Groupthink: Collective Delusions in Organizations and Markets. *Review of Economic Studies*, forthcoming.

Benabou, R., Tirole, J. (2002). Self-Confidence and Personal Motivation. *Quarterly Journal of Economics* 117 (3), 871-915.

Berlyne, D.E. (1954). A Theory of Human Curiosity. *British Journal of Psychology* 45 (3), 180-191.

Berlyne, D.E. (1957a). Conflict and Information-Theory Variables as Determinants of Human Perceptual Curiosity. *Journal of Experimental Psychology* 53, 399-404.

Berlyne, D.E. (1957b). Uncertainty and Conflict: A Point of Contact Between Information-Theory and Behavior-Theory Concepts. *The Psychological Review* 64 (6), 329-339.

Bernheim, B.D., Thomadsen, R. (2005). Memory and Anticipation. *Economic Journal* 115, 271-304.

Bhatia, S., Golman, R. (2013). Attention and Reference Dependence. Working Paper.

Bordalo, P., Gennaioli, N., Shleifer, A. (2012a). Salience Theory of Choice Under Risk. *Quarterly Journal of Economics* 127 (3), 1243-1285.

Bordalo, P., Gennaioli, N., Shleifer, A. (2012b). Salience in Experimental Tests of the Endowment Effect. *American Economic Review* 102 (3), 47-52.

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* 28, 95-160. New York: Academic Press.

Bromberg-Martin, E., Hikosaka, O. (2009). Midbrain Dopamine Neurons Signal Preference for Advance Information about Upcoming Rewards. *Neuron* 63, 119-126.

Brunnermeier, M., Parker., J. (2005). Optimal Expectations. *American Economic Review* 95 (4), 1092-1118.

Burks, S., Carpenter, J., Goette, L., Rustichini, A. (2013). Overconfidence and Social Signalling. *Review of Economic Studies* 80 (3), 949-983.

Cai, C.J., Meyer, R. (2013). Biases in Information Seeking to Avoid Catastrophic Financial Risk. Working Paper.

Caplan, L. (1995). Patient Delay in Seeking Help for Potential Breast Cancer. *Public Health Reviews* 23 (3), 263-274.

Caplin, A., Eliaz, K. (2003). AIDS Policy and Psychology: A Mechanism-Design Approach. *RAND Journal of Economics* 34 (4), 631-646.

Caplin, A., Leahy, J. (2001). Psychological Expected Utility Theory And Anticipatory Feelings. *Quarterly Journal of Economics* 116 (1) 55-79.

Carillo, J., Mariotti, T. (2000). Strategic Ignorance as a Self-Disciplining Device. *Review of Economic Studies* 67, 529- 544.

Cover, T., Thomas, J. (1991). *Elements of Information Theory*. New York: Wiley.

Dana, J., Weber, R., Kuang, J. (2007). Exploiting Moral Wiggle Room: Experiments Demonstrating an Illusory Preference for Fairness. *Economic Theory* 33 (1), 67-80.

Dekel, E., Lipman, B., Rustichini, A. (1998). Standard State-Space Models Preclude Unawareness. *Econometrica* 66 (1), 159-173.

DellaVigna, S., Pollet, J. (2009). Investor Inattention and Friday Earnings Announcements. *The Journal of Finance* 64 (2), 709-749.

Dillenberger, D. (2010). Preferences for One-Shot Resolution of Uncertainty and Allais-Type Behavior. *Econometrica* 78 (6), 1973-2004.

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., Schotter, A. (2010). Paying for Confidence: An Experimental Study of the Demand for Non-Instrumental Information. *Games and Economic Behavior* 70, 304-324.

Eliaz, K., Spiegler, R. (2006). Can Anticipatory Feelings Explain Anomalous Choices of Information Sources? *Games and Economic Behavior* 56 (1), 87-104.

Ellingsen, T., Johannesson, M. (2008). Anticipated Verbal Feedback Induces Altruistic Behavior. *Evolution and Human Behavior* 29 (2), 100-105.

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

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 (4), 3-30.

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

Gabaix, X., Laibson, D. (2006). Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets. *Quarterly Journal of Economics* 121, 505-540.

Gabaix, X., Laibson, D., Moloche, G., Weinberg, S. (2006). Costly Information Acquisition: Experimental Analysis of a Boundedly Rational Model. *American Economic Review* 96 (4), 1043-1068.

Galai, D., Sade, O. (2006). The 'Ostrich Effect' and the Relationship Between the Liquidity and the Yields of Financial Assets. *Journal of Business* 79, 2741-2759.

Ganguly, A., Tasoff, J. (2014). Fantasy and Dread: An Experimental Test of Attentional Anticipatory Utility.

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

Golman, R., Loewenstein, G. (2014). 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.

Hendry, D. (1969). Reinforcing Value of Information: Fixed-Ratio Schedules. In *Conditioned Reinforcement*, D. Hendry (Ed.), Homewood, IL: Dorsey, 300-341.

Heifetz, A., Meier, M., Schipper, B. (2006). Interactive Unawareness. *Journal of Economic Theory* 130 (1), 78-94.

Hirsh, J., Inzlicht, M. (2008). The Devil You Know: Neuroticism Predicts Neural Response to Uncertainty. *Psychological Science* 19 (10), 962-967.

Hirshleifer, J., Riley, J. (1979). The Analytics of Uncertainty and Information – An Expository Survey. *Journal of Economic Literature* 17 (4), 1375-1421.

Itti, L., Baldi, P. (2009). Bayesian Surprise Attracts Human Attention. *Vision Research* 49 (10), 1295-1306.

Janis, I. (1958). *Psychological Stress: Psychoanalytic and Behavioral Studies of Surgical Patients*. New York: Wiley.

Jenkins, H., Boakes, R. (1973). Observing Stimulus Sources that Signal Food or No Food. *Journal of the Experimental Analysis of Behavior* 20, 197-207.

Kadane, J., Schervish, M., Seidenfeld, T. (2008). Is Ignorance Bliss? *Journal of Philosophy* 105 (1), 5-36.

Kahneman, D. (1973). *Attention and Effort*. Englewood Cliffs, NJ: Prentice Hall.

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 (8), 963-973.

Kaplan, S. (1991). Beyond rationality: Clarity-based decision making. In *Environment, Cognition and Action: An integrative multidisciplinary approach*, T. Gärling and G. W. Evans (Eds.), New York: Oxford University Press

Kaplan, S. (1992). Environmental preference in a knowledge-seeking, knowledge-using organism. In *The Adapted Mind*, J. H. Barkow, L. Cosmides, and J. Tooby (Eds.), New York: Oxford University Press.

Karlsson, N., Loewenstein, G., McCafferty, J. (2004). The Economics of Meaning. *Nordic Journal of Political Economy* 30 (1), 61-75.

Karlsson, N., Loewenstein, G., Seppi, D. (2009). The Ostrich Effect: Selective Attention to Information. *Journal of Risk and Uncertainty* 38 (2), 95-115.

Karni, Edi, and Marie-Louise Vierø. (2013). "Reverse Bayesianism": A Choice-Based Theory of Growing Awareness. *American Economic Review* 103 (7), 2790-2810.

Köszegi, B. (2003). Health Anxiety and Patient Behavior. *Journal of Health Economics* 22 (6), 1073-1084.

Köszegi, B. (2006). Ego Utility, Overconfidence, and Task Choice. *Journal of the European Economic Association* 4 (4),673-707.

Köszegi, B. (2010). Utility from Anticipation and Personal Equilibrium. *Economic Theory* 44 (3), 415-444.

Köszegi, B., Szeidl, A. (2013). A Model of Focusing in Economic Choice. *Quarterly Journal of Economics* 128 (1), 53-104.

Kreps, D., Porteus, E. (1978). Temporal Resolution of Uncertainty and Dynamic Choice Theory. *Econometrica* 46 (1), 185-200.

Lerman, C., Hughes, C., Trock, B.J., Myers, R.E., Main, D., Bonney, A., Abbaszadegan, M.R., Harty, A.E., Franklin, B.A., Lynch, J.F., Lynch, H.T. (1999). Genetic Testing in Families with Hereditary Nonpolyposis Colon Cancer. *Journal of the American Medical Association* 281 (17), 1618-1622.

Lerman, C., Narod, S., Schulman, K., Hughes, C., Gomez-caminero, A., Bonney, G., Gold, K., Trock, B., Main, D., Lynch, J., Fulmore, C., Snyder, C., Lemon, S.J., Conway, T., Tonin, P., Lenoir, G., Lynch, H. (1996). BRCA1 Testing in Families with Hereditary Breast-Ovarian Cancer: A prospective study of patient decision making and outcomes. *Journal of the American Medical Association* 275 (24), 1885-1892.

Li, J. (2008). A Note on Unawareness and Zero Probability. PIER Working Paper No. 08-022.

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 (1), 20-42.

Locke, E., Latham, G. (1990). *A Theory of Goal Setting and Task Performance*. Englewood Cliffs, NJ: Prentice Hall.

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 (1), 75-98.

Loewenstein, G. (1999). Because it is there: The challenge of mountaineering... for utility theory. *Kyklos* 52, 315-344.

Lyter, D., Valdiserri, R., Kingsley, L, Amoroso, W., Rinaldo, C. (1987). The HIV Antibody Test: Why Gay and Bisexual Men Want or Do Not Want To Know Their Results. *Public Health Reports* 102, 468-474.

Modica, S., Rustichini, A. (1994). Awareness and Partitional Information Structures. *Theory and Decision* 37 (1), 107-124.

Modica, S., Rustichini, A. (1999). Unawareness and Partitional Information Structures. *Games and Economic Behavior* 27 (2), 265-298.

Meechan, G., Collins, J., Petrie, K. (2002). Delay in Seeking Medical Care for Self-Detected Breast Symptoms in New Zealand Women. *The New Zealand Medical Journal* 115, 257.

Miller, G., Galanter, E., Pribram, K.H. (1960). *Plans and the Structure of Behavior*. New York: Holt.

Möbius, M., Niederle, M., Niehaus, P., Rosenblat, T. (2011). Managing Self-Confidence: Theory and Experimental Evidence. NBER Paper No. w17014.

Nielsen, L. (1984). Unbounded Expected Utility and Continuity. *Mathematical Social Sciences* 8, 201-216.

Oster, E., Shoulson, I., Dorsey, E.R. (2013). Optimal Expectations and Limited Medical Testing: Evidence from Huntington Disease. *American Economic Review* 103, 804-30.

Picone, G., Sloan, F., Taylor, D. (2004). Effects of Risk and Time Preference and Expected Longevity on Demand for Medical Tests. *Journal of Risk and Uncertainty* 28, 39-53.

Powdthavee, N., Riyanto, Y. (2014). Would You Pay For Transparently Useless Advice? A Test of Boundaries of Beliefs in the Folly of Predictions. *Review of Economics and Statistics*, forthcoming.

Prokasy, W. (1956). The Acquisition of Observing Responses in the Absence of Differential External Reinforcement. *Journal of Comparative and Physiological Psychology* 49 (2), 131-134.

Rabin, M. (2000). Risk Aversion and Expected-utility Theory: A Calibration Theorem. *Econometrica* 68 (5), 1281-1292.

Ritov, I., Baron, J. (1990). Reluctance to vaccinate: Omission bias and ambiguity. *Journal of Behavioral Decision Making* 3 (4), 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 (4), 929-940.

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

Schelling, T. (1987). The Mind as a Consuming Organ. In *The Multiple Self*, J. Elster (Ed.), Cambridge: Cambridge University Press, 177-196.

Shannon, C. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal* 27, 379-423, 623-656.

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 (6), 787-791.

Stigler, G. (1961). The Economics of Information. *The Journal of Political Economy* 69 (3), 213-225.

Sweeny, K., Melnyk, D., Miller, W., Shepperd, J. (2010). Information Avoidance: Who, What, When, and Why. *Review of General Psychology* 14, 340-353.

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 (5), 1829-1863.

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.

von Neumann, J., Morgenstern, O. (1944). *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.

Wakker, P. (1988). Nonexpected Utility as Aversion of Information. *Journal of Behavioral Decision Making* 1 (3), 169-175.

Weaver, R., Frederick, S. (2012). A Reference Price Theory of the Endowment Effect. *Journal of Marketing Research*, forthcoming.

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 (1), 5-21.

Yariv, L. (2001). Believe and Let Believe: Axiomatic Foundations for Belief Dependent Utility Functionals. Cowles Foundation Discussion Paper No. 1344.