

Curiosity, Information Gaps, and the Utility of Knowledge

Russell Golman and George Loewenstein ^{*†}

April 16, 2015

Abstract

We use an information-gap framework to capture 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

Standard accounts of utility maximization tend to treat information as subservient to consumption. 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 of consumption. 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 anticipation (Loewenstein, 1987) apart from any material benefits the information might confer. At the same time, there are many situations in which people actively *resist* acquiring information (Sweeny et al., 2010; Golman et al., 2015). 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). The purpose of this paper is to enumerate the different motives underlying the acquisition or avoidance of information and show that they can be assimilated into Golman and Loewenstein's (2015a) 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. We generate predictions that are consistent with well-documented patterns of behavior as well as a number of novel predictions that are testable, but as yet untested.

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. Previous treatments of information avoidance in economics have generally derived information-avoidance from exotic 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.

We identify three motives underlying the desire to obtain or avoid information. As in the standard economic account, we recognize that information derives value by enabling one to make better subsequent decisions leading to superior expected outcomes. Secondly, individuals may seek or avoid information because they anticipate that what they discover will be pleasurable or painful. 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. We assume 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

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

information about questions they do not like thinking about. Lastly, people may want information to satisfy curiosity. We conceive of curiosity as a fundamental desire to fill ‘information gaps’ – specific uncertainties that one recognizes and is aware of (Loewenstein, 1994).

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 3.4 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 3.5 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

2 Theoretical Framework

2.1 Utility over Cognitive States

Following Golman and Loewenstein (2015a), we represent a person’s state of awareness with a set of activated questions $\mathcal{Q} = \{Q_1, \dots, Q_m\}$, where each question Q_i has a set of possible (mutually exclusive)

²Several 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).

answers $\mathcal{A}_i = \{A_i^1, A_i^2, \dots\}$.³ We let X denote a set of prizes. Denote the space of answer sets together with prizes as $\alpha = \mathcal{A}_1 \times \mathcal{A}_2 \times \dots \times \mathcal{A}_m \times X$. A cognitive state can then be defined by a 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 $\mathbf{w} = (w_1, \dots, w_m) \in \mathbb{R}_+^m$, as depicted in Table 1. A utility function is defined over cognitive states, written as $u(\pi, \mathbf{w})$.

Activated Questions	Possible Answers	Subjective Probabilities*	Attention Weights
Q_1	$\mathcal{A}_1 = \{A_1^1, A_1^2, \dots\}$	$[\pi_1(A_1^1), \pi_1(A_1^2), \dots]$	w_1
\vdots	\vdots	\vdots	\vdots
Q_m	$\mathcal{A}_m = \{A_m^1, A_m^2, \dots\}$	$[\pi_m(A_m^1), \pi_m(A_m^2), \dots]$	w_m
	Possible Prizes		
N/A	$X = \{x, x', x'', \dots\}$	$[\pi_X(x), \pi_X(x'), \dots]$	N/A

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

Table 1: Representation of a cognitive state.

The probability measure reflects a subjective probability judgment about the answers to the activated questions and the prizes that may be received. The subjective probability over these prizes is in general mutually dependent with the subjective probability over answers to activated questions. That is, material outcomes may correlate with answers about activated questions (and the answer to one question may correlate with the answer to another). We can consider a marginal distribution π_i that specifies the subjective probability of possible answers to question Q_i or π_X that specifies the subjective probability over prizes.⁴

The attention weights specify how much a person is thinking about each question and, in turn, how much the beliefs about those questions directly impact utility. The attention w_i on question Q_i is assumed to be strictly increasing in, and to have strictly increasing differences in, the question's importance γ_i and salience σ_i as well as to be strictly increasing in the surprise δ_i associated with it. We consider salience to be exogenous, whereas an individual's actions can endogenously affect importance and surprise. We return to characterize these concepts in Section 2.3 after first describing the choice setting.

A choice to acquire information is essentially a choice to accept a lottery over cognitive states because, ex ante, one cannot know what one will discover. 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 . Upon learning answer A_i to question Q_i , one's subjective probability measure over $\Delta(\alpha)$ changes from π^0 to $\pi^{A_i} = \pi^0(\cdot | A_i)$ due to conditioning of beliefs on the discovered answer, and one's attention weight vector changes from \mathbf{w}^0 to \mathbf{w}^{A_i} due to surprise. We assume Bayesian updating so that $\sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i) \pi^{A_i} = \pi^0$.⁵

³There is presumably an infinite set of latent questions that an individual could, in principle, ask, but is not currently aware of. We typically take the set of activated questions as given. Of course, in reality an individual's actions may lead to becoming aware of new questions. 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. Golman and Loewenstein (2015a) suggest that becoming aware of new questions is a source of wisdom which may be quite appealing to curious individuals if they anticipate that the possible answers are not worrisome.

⁴For any $\tilde{\mathcal{A}} \subseteq \mathcal{A}_i$, we have $\pi_i(\tilde{\mathcal{A}}) = \pi(\mathcal{A}_1 \times \dots \times \mathcal{A}_{i-1} \times \tilde{\mathcal{A}} \times \mathcal{A}_{i+1} \times \dots \times \mathcal{A}_m \times X)$.

⁵While 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.

Thus, acquiring information necessarily decreases expected entropy (see, e.g., Cover and Thomas, 1991, pg. 27).

Given that we think of acquiring information as accepting a lottery over cognitive states, standard von Neumann and Morgenstern (1944) assumptions imply an expected utility representation over cognitive states.⁶ 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 (π, \mathbf{w}) if $s^* \in \arg \max_{s \in \mathcal{S}} u(s \cdot (\pi, \mathbf{w}))$.⁷ We then define a utility function over cognitive states, contingent on the set of sequences of actions that may subsequently be chosen:

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

We also define the desirability of a sequence of actions s in cognitive state (π, \mathbf{w}) as the marginal utility relative to doing nothing:

$$D(s | \pi, \mathbf{w}) = u(s \cdot (\pi, \mathbf{w})) - u(\pi, \mathbf{w}).$$

2.2 The Desire for (or to Avoid) Information

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 (π^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

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

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.

⁶While 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 the 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 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). 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.

⁷We represent a sequence of contingent actions $s \in \mathcal{S}$ as a single operator on the prior cognitive state with the convention that each action operator passes through a distribution over cognitive states, akin to reduction of compound lotteries over cognitive states.

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

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

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

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

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

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

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

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

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

We address each of these motives, along with the phenomena they produce, in Section 3.

2.3 Attention

We assume that three factors – salience, importance, and surprise – contribute to the attention weights. Salience reflects the degree to which a particular context highlights the question. 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. We assume that a question Q_i is activated if and only if it has positive salience $\sigma_i > 0$.

A question is important to the extent that one’s utility depends on the answer. To characterize the importance of question Q_i , we consider the probabilities of discovering any possible answer $A_i \in \mathcal{A}_i$ (or, omitting answers thought to be impossible, in the support of the individual’s belief about the question, $\text{supp}(\pi_i)$) and the utilities of the cognitive states $(\pi^{A_i}, \mathbf{w}^{A_i})$ that would result from discovering each possible answer A_i contingent upon having the set \mathcal{S} of sequences of actions subsequently available. We assume that the importance γ_i of question Q_i is a function of the subjective distribution of utilities that would result from different

answers to the question,

$$\gamma_i = \phi \left(\left\langle \pi_i^0(A_i), U(\pi^{A_i}, \mathbf{w}^{A_i} | \mathcal{S}) \right\rangle_{A_i \in \text{supp}(\pi_i^0)} \right), \quad (6)$$

that increases with mean-preserving spreads of the distribution of utilities and that is invariant with respect to constant shifts of utility.⁸ According to our definition, importance depends on utility, which in turn depends on the attention weight, but importance also contributes to attention weight. This circularity captures the dynamic processes giving rise to obsession: attention to a question raises its importance, and the elevated importance gives rise to intensified attention. Nevertheless, raising the stakes associated with a question clearly increases its importance. Yet, if an answer is known with certainty, then by our definition nothing is at stake, so the underlying question is no longer important. We assume that there is some delay after acquiring information before an individual can adapt to it and adjust the importance of the questions being addressed.

The surprise one experiences upon acquiring new information reflects the degree to which this information changes existing beliefs. We assume that when the answer A_j (to question Q_j) is learned (thereby providing information about associated questions and causing their subjective probabilities to be updated), the degree of surprise associated with revised 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 to new information it ceases to be surprising.

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. The *belief resolution effect* refers to the resulting dynamic pattern of attention. When an answer is learned with certainty, there is an immediate boost in attention weight on it, but as the person adapts over time, this attention weight falls to a lower level. Janis (1958) implicitly recognized the belief resolution effect, observing that surgical patients getting information about their upcoming procedures initially worry more about the surgery but subsequently experience less anxiety.

⁸This 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.

2.4 Example: Contemplating Teaching Ratings

To illustrate the modeling of a cognitive state, consider a college professor deciding whether or not to look at her teaching ratings. The set of activated questions (and possible answers) might include: “How many of my students liked my teaching?” (0, 1, 2, . . .); “Did they applaud on the last day of class?” (yes/no); “How good a teacher am I?” (great, good, so-so, bad, awful); “Will I get tenure?” (yes/no). Prior belief about the first question might be quite uncertain. The answer to the second question, on the other hand, might already be known with certainty. There may or may not be much uncertainty about the third and fourth questions. All of these beliefs (to the extent they are uncertain) are jointly dependent. 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.

We may not know the precise levels of attention paid to each of the questions above, but we can make some comparisons across different contexts. When the professor first becomes aware that the teaching ratings have been compiled, the question of how many students liked her teaching becomes somewhat salient. If colleagues proceed to discuss their ratings at a faculty meeting, the question then becomes more salient. On the other hand, if the professor is presently giving a research talk, the question becomes much less salient. If the professor teaches at a liberal arts college, her teaching ratings should be more correlated with her prospects for tenure than if she teaches at a research university. Thus, we would expect the question of how many students liked the professor’s teaching to be more important if she is at the liberal arts college than if she is at the research university. Similarly, if she cares more about getting tenure, e.g., if she is committed to an academic career as opposed to considering a private job offer, these questions would be more important to her. Lastly, if she has recently discovered some criticism of her teaching, the surprise associated with her shifted expectations would cause her to focus more attention on the question of her teaching ratings, especially if the criticism was unexpected. But, if she has gotten used to such criticism and become convinced she is a bad teacher, the surprise will have worn off and it will cease to seem so important and she may pay little attention to her teaching ratings, as the belief resolution effect occurs.

The possibility of enrolling in a teacher improvement class would be relevant for the decision whether to look at the teaching ratings. Given the option subsequently to enroll in the teacher improvement class or not to enroll in the class, we model looking at the ratings as resolving a lottery over cognitive states, each of which having utility that is conditional on making the optimal choice of the subsequent action (to enroll or not to enroll). Finding out the teaching ratings would inform the decision whether or not to enroll in the class, so the information has instrumental value. The resolution into a specific cognitive state still impacts utility directly as well. If the professor comes to believe she is a good (bad) teacher and to pay more attention to this belief, she will feel good (bad). Motivated attention supports looking (not looking) if she anticipates good (bad) news. On the other hand, if the professor avoids looking at her teaching ratings entirely, the uncertainty about them may continue to bother her. Curiosity pushes her to resolve this uncertainty.

2.5 A Specific Utility Function

To make precise predictions about preferences for (or to avoid) information, we consider a specific utility function. This utility function should allow attention weights to modulate the impact of beliefs on utility, should allow that certain beliefs may have intrinsic value or valence,⁹ and should capture a general aversion to uncertainty or preference for clarity. The desire for clarity is central to the phenomenon of curiosity, which we discuss in greater depth in Section 3.2. 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 3.3.

To capture a preference for clarity, we make use of a common measure of the uncertainty about a particular question Q_i : the entropy of the subjective probability distribution over its answers, $H(\pi_i) = -\sum_{A_i \in \mathcal{A}_i} \pi_i(A_i) \log \pi_i(A_i)$ with the convention that $0 \log 0 = 0$ (Shannon 1948).¹⁰ We consider the utility function

$$u(\pi, \mathbf{w}) = \sum_{x \in X} \pi_X(x) v_X(x) + \sum_{i=1}^m w_i \left(\sum_{A_i \in \mathcal{A}_i} \pi_i(A_i) v_i(A_i) - H(\pi_i) \right). \quad (7)$$

The first term describes expected utility over prizes and the remaining terms describe the utilities of beliefs about each activated question, amplified by the attention weights on each of these questions. Golman and Loewenstein (2015a) describe properties that characterize this utility function. Our primary focus here is exploring the patterns of behavior that this utility function implies.

3 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 (7), 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.

3.1 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

⁹We can identify as positive (neutral / negative) beliefs those for which increasing attention on the belief increases (does not affect / decreases) utility.

¹⁰The base of the logarithm in the entropy formula is arbitrary and amounts to a normalization parameter.

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 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).

3.2 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.¹¹ 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

¹¹Recent 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.

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.

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 2) 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 2: 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 3.3, 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 (7), 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.¹² (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

¹²Entropy times attention weight, as built into Equation (7), 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).

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 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 2.3. 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.

3.3 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 2.3) 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.

3.4 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 (7), 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}_ν^* 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 (7) 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.¹³ For example, willingness to pay for an

¹³In 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

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 (Meehan 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 (7)) that

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.

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).

Falk and Zimmermann (2014) designed a clever laboratory experiment to test our prediction. Subjects in two conditions chose whether to obtain costless information about whether they would receive fifteen minutes later a series of electric shocks – anticipation of the shocks clearly being a negative belief. In one condition subjects could distract themselves by playing a quiz game, thus decreasing the salience of the question of whether they would receive the shocks. In the control condition no distraction was available, so the possibility of impending shocks remained highly salient. Consistent with our theory, when the possibility of receiving electric shocks was more salient subjects were much more likely to obtain the information. Similarly, we might hypothetically 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 and especially if the patient was meeting with the doctor with no opportunity to find a distraction, the patient might then prefer to be informed of the test results.¹⁴

3.5 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

¹⁴If 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.

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.

4 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.

One fundamental assumption of our model is that uncertainty has intrinsic costs (leading to the desire for clarity). Flipped around, we could say that knowledge has intrinsic value (Karlsson et al., 2004). Conventional economic theory can certainly accommodate choices to devote significant time, money, and effort to developing knowledge that is unlikely to confer any material benefits. One could posit, for example, that the wine connoisseur develops expertise because it improves the consumption experience – i.e., that,

in effect, the connoisseur and the novice 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 knowledge, it seems much more parsimonious to accept knowledge as a direct source 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.

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.¹⁵
- H3 Individuals who discount the future less should be more likely to expose themselves to information relating to negative beliefs.¹⁶
- 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.

¹⁵The 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.

¹⁶While 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.

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.

In a companion paper (Golman and Loewenstein, 2015b) we argue that the information-gap concept applied here to information acquisition and avoidance also 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. 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.

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.

Appendix

Proof of Proposition 1

Conditions 1 and 2 imply that attention weight has been made stronger on some pairwise dependent question Q_{j^*} , $\hat{w}_{j^*}^0 > w_{j^*}^0$, and no weaker on other pairwise dependent questions Q_j , $\hat{w}_j^0 \geq w_j^0$. 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 (see Golman and Loewenstein, 2015a). 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

$$\hat{D}_i - D_i = \sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i) (\hat{u}(\pi^{A_i}, \hat{\mathbf{w}}^{A_i}) - u(\pi^{A_i}, \mathbf{w}^{A_i})) - (\hat{u}(\pi^0, \hat{\mathbf{w}}^0) - u(\pi^0, \mathbf{w}^0)), \quad (8)$$

where the terms representing instrumental value have vanished by assumption. We expand the utility functions according to Equation (7) and group terms in a utility difference

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

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 $\hat{w}_j^{A_i} - \hat{w}_j^0 = w_j^{A_i} - w_j^0$ or equivalently $\hat{w}_j^{A_i} - w_j^{A_i} = \hat{w}_j^0 - w_j^0$. This allows us to extract a common factor of $\hat{w}_j^0 - w_j^0$ in the last term of the expansion of Equation (8):

$$\begin{aligned} \hat{D}_i - D_i &= \left(\sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i) \pi_X^{A_i}(x^*) - \pi_X^0(x^*) \right) (\hat{v}_X(x^*) - v_X(x^*)) + \\ &\quad \left(\sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i) \hat{w}_\nu^{A_i} \pi_\nu^{A_i}(A_\nu^*) - \hat{w}_\nu^0 \pi_\nu^0(A_\nu^*) \right) (\hat{v}_\nu(A_\nu^*) - v_\nu(A_\nu^*)) + \\ &\quad \sum_{j=1}^m (\hat{w}_j^0 - w_j^0) \left(\sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i) \left(\sum_{A_j \in \mathcal{A}_j} (\pi_j^{A_i}(A_j) - \pi_j^0(A_j)) v_j(A_j) - H(\pi_j^{A_i}) + H(\pi_j^0) \right) \right). \end{aligned}$$

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

$$\begin{aligned} \hat{D}_i - D_i &= \sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i) \pi_\nu^{A_i}(A_\nu^*) (\hat{w}_\nu^{A_i} - \hat{w}_\nu^0) (\hat{v}_\nu(A_\nu^*) - v_\nu(A_\nu^*)) + \\ &\quad \sum_{j=1}^m (\hat{w}_j^0 - w_j^0) \left(H(\pi_j^0) - \sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i) H(\pi_j^{A_i}) \right). \end{aligned}$$

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 $\hat{w}_\nu^{A_i} - \hat{w}_\nu^0 \geq 0$, 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., $H(\pi_{j^*}^0) - \sum_{A_i} \pi_i^0(A_i)H(\pi_{j^*}^{A_i}) > 0$. With $\hat{w}_{j^*}^0 > w_{j^*}^0$ 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 $\pi \rightarrow \hat{\pi}$ specified by condition 4. Recognizing that $\hat{\pi}_j^0 = \pi_j^0$ for all j , we have $u(\hat{\pi}^0, \mathbf{w}^0) = u(\pi^0, \mathbf{w}^0)$ (again using Equation (7)), and the change in the desire for information simplifies as

$$\hat{D}_i - D_i = \sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i) (u(\hat{\pi}^{A_i}, \hat{\mathbf{w}}^{A_i}) - u(\pi^{A_i}, \mathbf{w}^{A_i})). \quad (9)$$

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

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

Plugging this into Equation (9) and simplifying with the law of total probability, we obtain

$$\hat{D}_i - D_i = \sum_{A_i \in \mathcal{A}_i} \pi_i^0(A_i) [(\hat{w}_\nu^{A_i} - w_\nu^{A_i}) v_\nu(\hat{\pi}_\nu^{A_i}) + w_\nu^{A_i} (H(\pi_\nu^{A_i}) - H(\hat{\pi}_\nu^{A_i}))] \quad (10)$$

We know $\hat{w}_\nu^{A_i} \geq w_\nu^{A_i}$ 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 (10) is nonnegative. When the questions are pairwise independent, conditioning on A_i does not change the belief about Q_ν , so $\pi_\nu^{A_i} = \pi_\nu^0 = \hat{\pi}_\nu^0$. Conditioning on the answer A_i strictly decreases (in expectation) the entropy of the belief about a pairwise dependent question, so

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

Thus, returning to Equation (10), 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.
- 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.
- 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.
- 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.
- 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.
- Ellingsen, T., Johannesson, M. (2008). Anticipated Verbal Feedback Induces Altruistic Behavior. *Evolution and Human Behavior* 29 (2), 100-105.
- Falk, A., Zimmermann, F. (2014). Beliefs and Utility: Experimental Evidence on Preferences for Information.
- 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

- 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. (2015a). An Information-Gap Framework for Capturing Preferences About Uncertainty. Proceedings of the Fifteenth Conference on Theoretical Aspects of Rationality and Knowledge.
- Golman, R., Loewenstein, G. (2015b). Information Gaps for Risk and Ambiguity. Working Paper.
- Golman, R., Hagmann, D., Loewenstein, G. (2015). Information Avoidance. In prep.
- 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.
- Hirshleifer, J., Riley, J. (1979). The Analytics of Uncertainty and Information – An Expository Survey. *Journal of Economic Literature* 17 (4), 1375-1421.
- 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. (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.
- 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.
- 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.

- 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.
- 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.
- 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.
- 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.
- Möbius, M., Niederle, M., Niehaus, P., Rosenblat, T. (2011). Managing Self-Confidence: Theory and Experimental Evidence. NBER Paper No. w17014.
- 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.
- Shannon, C. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal* 27, 379-423.
- 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.