informs
http://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
Articles in Advance, pp. 1–23
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

The Demand for, and Avoidance of, Information

Russell Golman,[a] George Loewenstein,[a] Andras Molnar,[b] Silvia Saccardo[a]

[a] Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213; [b] Booth School of Business, University of Chicago, Chicago, Illinois 60637

Contact: rgolman@andrew.cmu.edu, https://orcid.org/0000-0002-3414-223X (RG); gl20@andrew.cmu.edu, https://orcid.org/0000-0003-2790-0474 (GL); andras.molnar@chicagobooth.edu, https://orcid.org/0000-0002-4341-9659 (AM); ssaccard@andrew.cmu.edu, https://orcid.org/0000-0001-7560-8262 (SS)

Received: May 14, 2020
Revised: March 29, 2021; July 7, 2021
Accepted: August 15, 2021
Published Online in Articles in Advance: December 15, 2021

https://doi.org/10.1287/mnsc.2021.4244

Abstract. Management scientists recognize that decision making depends on the information people have but lack a unified behavioral theory of the demand for (and avoidance of) information. Drawing on an existing theoretical framework in which utility depends on beliefs and the attention paid to them, we develop and test a theory of the demand for information encompassing instrumental considerations, curiosity, and desire to direct attention to beliefs one feels good about. We decompose an individual's demand for information into the desire to refine beliefs, holding attention constant, and the desire to focus attention on anticipated beliefs, holding these beliefs constant. Because the utility of resolving uncertainty (i.e., refining beliefs) depends on the attention paid to it and more important or salient questions capture more attention, demand for information depends on the *importance* and *salience* of the question(s) it addresses. In addition, because getting new information focuses attention on one's beliefs and people want to savor good news and ignore bad news, the desire to obtain or avoid information depends on the *valence* (i.e., goodness or badness) of anticipated beliefs. Five experiments ($n = 2{,}361$) test and find support for these hypotheses, looking at neutrally valenced as well as ego-relevant information. People are indeed more inclined to acquire information (a) when it feels more important, even if it cannot aid decision making (Experiments 1A and 2A); (b) when a question is more salient, manipulated through time lag (Experiments 1B and 2B); and (c) when anticipated beliefs have higher valence (Experiment 2C).

History: Accepted by Yan Chen, behavioral economics and decision analysis.
Funding: The project was partially supported by the National Science Foundation [Grant SES-1919453 to S. Saccardo].
Supplemental Material: The data files and online appendix are available at https://doi.org/10.1287/mnsc.2021.4244.

Keywords: curiosity • information gap • motivated attention • ostrich effect

1. Introduction

Good decision making depends on the information people have, but people may be wary of information that challenges their existing beliefs, warns of impending bad outcomes, or addresses problems not currently on their radar. They may be more inclined to look at information that feels reassuring or that simply grabs their attention. When managers or policy makers want to disseminate information to other decision makers (e.g., customers, strategic partners, shareholders, etc.), they need to cut through the cacophony of competing information campaigns and in some cases, overcome avoidance of potentially unpleasant information to get people to listen. In these situations and many others, it would be helpful to know when and why people either seek out information or avoid it. We develop a unified theory offering predictions about when, and how strongly, people will want to acquire or avoid information, and we provide evidence from five experiments supporting three of our theory's predictions.

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. The theory predicts that (outside of strategic situations) valid information will never be valued negatively because at worst, it can be ignored (i.e., not taken into account in decision making). Yet, there are many situations in which people actively *resist* acquiring information (Hertwig and Engel 2016, Golman et al. 2017). 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 would provide valuable information for decision making (e.g., whether to obtain treatment). At the same time, people also seek out information, such as celebrity gossip, that does not benefit decision making (Kruger and Evans 2009, Eliaz and Schotter 2010, Hsee and Ruan 2016).

Our theory of information seeking and avoidance is unique in highlighting the role of attention to specific beliefs. It assumes that people derive utility from thinking about specific beliefs, so information not only informs decision making but also, directly impacts utility by refining beliefs and redirecting the focus of attention. Golman and Loewenstein (2018a) present the attention-weighted belief-based utility function and propose that attention depends on endogenous psychological constructs ("importance" and "surprise") along with exogenous conditions ("salience"). Here, we apply that framework to generate testable predictions about information seeking and avoidance, identifying specific contextual factors that stimulate demand for information. The theory reconciles previously disconnected sets of empirical findings across different domains and makes new predictions that we test, and find support for, in this paper. Although other theories recognize that people sometimes seek useless information or avoid useful information, our theory is unique in identifying contextual factors affecting attention that determine when these nonstandard informational preferences occur.

The theory incorporates two distinct motives underlying the desire to obtain or avoid information on top of the traditional instrumental value of information. First, individuals may seek or avoid information because they anticipate that what they discover will be pleasurable or painful (as in Caplin and Leahy 2001, Kőszegi 2010). Beliefs can be pleasurable or painful (i.e., positively or negatively *valenced*) for many reasons—for instance, they can evoke anticipatory emotions or affect one's self-concept. 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 nonzero expected impact on belief-based utility (see Eliaz and Spiegler 2006). However, we assume that obtaining news tends to increase attention to revised beliefs (as in Gabaix et al. 2006) (i.e., to know something, at least at the moment of finding out, has a greater impact on utility than to merely suspect it (see Karlsson et al. 2009)). This impact-magnifying effect of new information leads people to seek information about questions they like thinking about and avoid information about questions they do not like thinking about. Second, people may seek information to satisfy curiosity. There are countless things people want to know despite having no practical use for the information. People incur real costs to indulge their curiosity (Kruger and Evans 2009, Eliaz and Schotter 2010, Hsee and Ruan 2016, Alos-Ferrer et al. 2018). Curiosity may result from a fundamental drive to reduce uncertainty and thus, make sense of one's environment (see Gottlieb et al. 2013, Kidd and Hayden 2015, Buyalskaya and Camerer 2020). Yet, people become curious about specific questions, and not others. We conceive of curiosity as a basic urge to resolve uncertainty about specific questions that capture attention—*information gaps* (Loewenstein 1994). All in all, our theory asserts that information demand is driven by instrumental value, intrinsic desire for reducing uncertainty about salient questions (curiosity), and motivated attention to good versus bad news.[1] Curiosity may inspire people to acquire noninstrumental information, whereas the motive to direct attention to more favorable beliefs may inspire people to avoid potentially useful information.

To illustrate the trade-offs involved in information acquisition or avoidance, consider a person deciding whether to obtain some performance feedback: for example, an employee deciding whether to read her manager's evaluation of her. Getting the feedback would inform a decision about how to improve in the future, so the information has instrumental value. The information would impact hedonic utility as well. If the employee learns that her manager is satisfied with her performance and pays more attention to this belief, she will feel good; if she learns of a poor evaluation and dwells on it, she will feel bad. A desire to focus attention on, or away from, this ego-relevant belief promotes looking if she anticipates good news and not looking if she anticipates bad news. Yet, if she remains uncertain and cannot forget about it, nagging curiosity may push her to find out.

We analyze how attentional factors affect curiosity and how the valence of beliefs affects the motive to focus on or ignore these beliefs, and we thus derive testable predictions about when people will be more motivated to seek or avoid information. The theory accounts for a range of empirical findings that had not yet been fit together within a coherent, comprehensive model and also generates new hypotheses about how attention affects demand for information. It predicts that making a question more important stimulates curiosity for finding out the answer because higher importance (defined as greater potential impact on utility, not necessarily corresponding to greater instrumental usefulness) directs more attention to the presence of an information gap. Similarly, it predicts that making a question more salient (defined as directing attention to it through contextual factors) also stimulates curiosity for the answer. People may seek this information even for epistemic questions for which they have no strict preferences between the possible answers. Additionally, when beliefs have valence, the motive to direct attention away from questions one does not like thinking about, and toward questions one likes thinking about, leads to stronger desire for information when the anticipated answers have higher valence.

We test these predictions in five online experiments ($n = 2{,}361$) that focus on the demand for noninstrumental information. We conduct these experiments in two domains: (1) offering neutrally valenced, epistemic information (i.e., individuals can find out the answer to a riddle) to capture the effect of attention on curiosity *absent* any desire to redirect attention away from bad news and (2) offering ego-relevant information, which could be construed as good or bad news (i.e., individuals can learn about their performance on a test). In both domains, we show that, absent any instrumental value, amplifying attention by manipulating the perceived importance of a given question (an answer to a riddle in Experiment 1A or one's performance on a test in Experiment 2A) stimulates the desire to fill the information gap (i.e., to find out the answer to the riddle or to reveal the results of the test). In both domains, we also show that demand for information is higher when the information gap had been opened more recently (i.e., when it is more salient (Experiments 1B and 2B)). Finally, in the ego-relevant domain, where the motive to redirect attention away from bad news is at play, we manipulate the valence of feedback the participants could receive about their performance on a test and confirm the model's prediction that demand for noninstrumental information is higher when individuals expect to receive good news (i.e., when the anticipated answers have higher valence). This set of findings cannot be explained by other theories of information acquisition or avoidance.

Our theory coherently brings together demand for noninstrumental information and avoidance of potentially useful information, accounting for specific patterns of each that have not been jointly reconciled until now and capturing the role of contextual factors that have been overlooked in existing work. Previous treatments of the demand for noninstrumental information have posited that people have intrinsic preferences regarding the resolution of uncertainty (Grant et al. 1998, Cabrales et al. 2013, Ely et al. 2015) or intrinsic preferences for information that may give them confidence in decisions they are about to make (Asch et al. 1990, Eliaz and Schotter 2010). The predictions of these models do not distinguish different sources of uncertainty with identical probabilistic structure. For example, if two football games went into overtime with identical probabilities of one's preferred team winning and identical probabilities of all interim events, these models would predict the same level of curiosity to see each game play out, even if one game was shown live and the other on tape delay or even if the viewer had just tuned into one game but had attentively watched the other game from the beginning. By contrast, we distinguish between different information gaps according to the attention devoted to each, so that a person may be very curious to fill one gap while being indifferent to another with similar probabilistic structure. We thus predict that curiosity will vary with situational determinants (see Loewenstein 1994) because a variety of contextual factors can affect attention to information gaps and thus, demand for information. Our experiments provide supportive evidence that amplifying attention by increasing the importance or the salience of an information gap strengthens the preference for filling it.

Previous treatments of information avoidance have generally derived it (a) from nonstandard risk preferences (e.g., Kreps and Porteus 1978, Wakker 1988, Grant et al. 1998, Dillenberger 2010, Andries and Haddad 2020), (b) from belief-based utility with risk aversion or loss aversion (e.g., Caplin and Leahy 2001; Kőszegi 2006, 2010; Karlsson et al. 2009; Pagel 2018) (i.e., from assuming that negative surprises have more impact than positive surprises), or (c) from optimism (i.e., from assuming that people can choose favorable beliefs in the absence of information) (e.g., Brunnermeier and Parker 2005, Oster et al. 2013). A limitation of many of these models is that they make the unrealistic prediction that a person who avoids information when anticipating bad outcomes must also avoid information when anticipating good outcomes, as highlighted by Eliaz and Spiegler (2006).[2] Our theory is not subject to this critique because it accounts for information avoidance as a result of the desire to avoid increasing attention to unpleasant beliefs. Thus, avoiding information when anticipating bad outcomes is consistent with demanding information when anticipating good outcomes, in line with the results of our third experiment as well as previous empirical findings (Karlsson et al. 2009, Eil and Rao 2011, Ganguly and Tasoff 2016, Gigerenzer and Garcia-Retamero 2017). Furthermore, models based on anticipatory utility or on optimism also predict no preference for information when people do not care at all about the outcomes (e.g., for epistemic information, such as finding out the answer to a riddle). By contrast, our theory not only accounts for the finding that information avoidance is more common when beliefs are more negatively valenced but also predicts that people experience curiosity for the answer to a question even when all answers have neutral valence, simply to resolve uncertainty about a question on one's mind.

The remainder of the paper is organized as follows. Section 2 reviews the theoretical framework introduced by Golman and Loewenstein (2018a), develops its application to preferences for information more fully than in that paper, and offers testable predictions about the demand for information. Section 3 presents experiments testing these predictions, demonstrating that different factors that govern attention to an information gap or that influence the valence of potential beliefs affect demand for noninstrumental information. Section 4 discusses additional predictions of the theory, including other drivers of curiosity and implications for information avoidance and individual welfare. Section 5 concludes.

2. Theory

2.1. Attention-Based Utility

Our theory incorporates a form of belief-based utility in which attention to beliefs modulates their impact on utility. We represent attention to different beliefs using the Golman and Loewenstein (2018a) question and answer framework, presented in Online Appendix A. In this framework, we define an information gap as a question that one is aware of but for which one is uncertain between possible answers. We thus distinguish the specific uncertainties that a person is paying attention to from the many other things the person does not know and does not think about. Utility depends on beliefs and the attention paid to them but not on uncertainties that do not capture attention. We denote the utility function as $u(\pi, \mathbf{w})$, where π is a probability measure representing beliefs and $\mathbf{w}$ is a vector representing attention to each question (i.e., each belief). To make precise predictions, we use the specific utility function proposed by Golman and Loewenstein (2018a) and presented in Online Appendix A. It assumes that beliefs contribute to utility to the extent a person pays attention to them, allows certain beliefs to have intrinsic value or valence, and captures a general aversion toward uncertainty, the latter of which is also evident from typical patterns of neural activity (see, e.g., Hirsh and Inzlicht 2008, Gottlieb et al. 2014). We do not treat beliefs (or attention) as choice variables—something that an individual could freely choose—but focus on decisions of whether to acquire information to influence beliefs (and attention).

2.2. Attention

Golman and Loewenstein (2018a) propose that the importance and salience of a question affect the attention devoted to it, as does the surprise associated with any discovery about it.

A question is important to the extent that one's utility depends on the answer.[3] Notably, questions can be important without having instrumental value, as a person may care about the answer even if it does not affect decision making. We characterize the *importance* of a question as a function of the distribution of utilities that would result from different answers to the question. If this distribution becomes more (or less) spread out, the question becomes more (or less) important. For instance, an opportunity to gain or lose a large amount of money (or self-esteem or hope) depending on the answer to a question can make that question more important, even if knowing the answer has no instrumental value (as in Experiments 1A and 2A in Section 3). If an answer is known with certainty, then by our definition, there is no spread in possible utilities, so the underlying question is no longer important. However, we assume that acquiring information, and revising beliefs, does not affect the importance of the questions being addressed until the person adapts to the new beliefs.

Salience reflects the degree to which a particular context highlights a question, possibly because of the passage of time (as in Experiments 1B and 2B in Section 3), the presence of distractions, comparison and contrast, or social cues. For example, awareness that performance feedback has been provided to a peer would make questions about one's own performance on the same task more salient. These questions would be more salient immediately after completing the task than they would be days later. Also, working on an engaging, unrelated task would make these questions less salient.

The *surprise* one experiences upon acquiring new information reflects the degree to which this information changes existing beliefs. We assume that the degree of surprise associated with a revised belief about a question when some information is obtained is the Kullback–Leibler divergence of the revised belief about that question against the prior belief about that question. 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.

2.3. Preferences About Information

A choice to acquire information is essentially a choice to accept a lottery over beliefs (and attention) because ex ante, one cannot know what one will discover. In the absence of information, an individual has beliefs π^0 and attention $\mathbf{w}^0$. Upon learning answer A_i to question Q_i, beliefs change from π^0 to $\pi^{A_i} = \pi^0(\cdot \mid A_i)$ because of conditioning of beliefs on the discovered answer, and attention changes from $\mathbf{w}^0$ to $\mathbf{w}^{A_i}$ because of surprise. We assume *Bayesian updating* as well as an *expected utility* representation for the utility of a lottery over beliefs and attention.[4] Assuming that backward induction is used to evaluate a sequence of actions, where early actions may reveal information that will inform later actions, we define a utility function contingent[5] on the set $\mathcal{S}$ of sequences of actions that may subsequently be chosen:

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

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.

The desire for information answering question Q_i, given prior belief π^0 and attention $\mathbf{w}^0$ and with a set $\mathcal{S}$ of subsequent sequences of actions available, is

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

Naturally, when D_i is positive (or negative), an individual seeks (or avoids) 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 ($\pi^0 \to \pi^{A_i}$).
3. The information may change the attention weights ($\mathbf{w}^0 \to \mathbf{w}^{A_i}$).

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 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 (Hirshleifer and Riley 1979). The instrumental value of information answering question Q_i, when the set $\mathcal{S}$ of subsequent sequences of actions is available, is

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

where $D(s \mid \pi, \mathbf{w}) = u(s \cdot (\pi, \mathbf{w})) - u(\pi, \mathbf{w})$ is the desirability of a sequence of actions s relative to doing nothing. In our framework, information can have instrumental value either if it supports a better choice among subsequent actions or if it makes an intended subsequent action more (or less) attractive.[6] As an example of this latter form of instrumental value, 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). Similarly, a dieter might refuse to read nutritional facts about a dessert he has already decided to eat so he can enjoy it unencumbered by thoughts of its health consequences or to avoid subsequent guilt (Woolley and Risen 2018). Although instrumental value derived from the usefulness of information is positive whenever dynamic consistency holds, it may be negative if dynamic consistency is violated, possibly because of moral wiggle room (Dana et al. 2007), temptation (Woolley and Risen 2018), motivation maintenance (Bénabou and Tirole 2002), or the curse of knowledge (Camerer et al. 1989). Additionally, instrumental value derived from complementarity or substitutability with subsequent actions can be positive or negative.

Curiosity is an intrinsic desire for knowledge that occurs when an individual becomes aware of a gap in his or her knowledge that could potentially be filled by information (Loewenstein 1994). In our framework, an information gap opens when a specific unanswered question Q_i captures attention. We identify curiosity for the answer to question Q_i as

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

This is the gain in utility from updating beliefs, holding attention fixed. In general, D_i^{C} could be positive or negative, but if we apply the Golman and Loewenstein (2018a) utility function with a cost of uncertainty (in the form of entropy times attention weight), then this term simplifies to an attention-weighted expected reduction in entropy because ex post beliefs cannot be expected to be any better or worse than ex ante beliefs. This aligns with our conception of curiosity as an urge to resolve uncertainty about the questions that capture attention. With this utility function, $D_i^{\mathrm{C}} \geq 0$ because acquiring information decreases expected entropy (see Cover and Thomas 1991, p. 27). Thus, as a representation of curiosity, it acts exclusively as a motive for information seeking.

Consistent with the view that curiosity supports sensemaking (Chater and Loewenstein 2015, Liquin and Lombrozo 2020) rather than simple uncertainty reduction (Cabrales et al. 2013), our theory predicts that curiosity arises only for information that addresses one or more questions that a person is already asking (see Berlyne 1954). 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, a well-known, reliable measure of attention (Kahneman 1973).

Motivated attention to (or away from) anticipated beliefs arises from the impact of obtaining information on attention. We express this as

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

Motivated attention may either contribute to the desire to seek information or drive avoidance of information, depending on the valence of anticipated beliefs. According to the theory, revising a belief attracts attention through surprise. Naturally, people prefer to think about positive rather than negative situations, so they tend to desire information about questions with positively valenced answers and to avoid information about questions with negatively valenced answers. For example, most people enjoy opening a gift (in addition to receiving it) because they experience a pleasant surprise. On the other hand, most people do not enjoy going to see the doctor for a diagnosis.

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

Theorem 1.

$$D_i = D_i^{\mathrm{IV}} + D_i^{\mathrm{C}} + D_i^{\mathrm{MA}}. \tag{6}$$

Theorem 1 states that three motives contribute to the desire for information: instrumental value, curiosity, and motivated attention.

2.4. Predictions About the Demand for Information

Analyzing each of these motives with the attention-based utility function yields a number of new testable predictions, and we focus especially on predictions driven by the role of attention. First, observe that curiosity, as defined in Equation (4), depends on the prior attention given to an information gap. This level of attention, based on the salience and importance of the information gap, modulates the curiosity that arises from a potential reduction in uncertainty. Next, observe that motivated attention, as defined in Equation (5), depends on the change in attention upon obtaining information. This change in attention, based on surprise, interacts with the valence of anticipated beliefs. So, according to our theory, directing attention to (or away from) the presence of an information gap can increase (or decrease) demand for information because of the curiosity motive, and shifting the valence of potential beliefs can affect demand for information because of the motive to direct attention away from bad news (and toward good news). This leads to our primary hypotheses about the effects of importance, salience, and valence on the demand for information, presented here, as well as additional predictions that we discuss in our discussion section.[7]

Hypothesis 1. *Increasing the importance of a question increases demand for information pertaining to that question.*

Hypothesis 2. *Increasing the salience of a question increases demand for information pertaining to that question.*

Hypothesis 3. *Uniformly increasing the valence of the answers to a question increases demand for information pertaining to that question.*

Hypothesis 1 holds because curiosity tends to be stronger about questions that an individual considers as more *important*. Although information with higher instrumental value is typically seen as more important, recall that, by our definition, having instrumental value is not a necessary condition for importance. People may perceive information as important even when it cannot affect their future decisions simply because it does affect their utility. As people tend to care about material outcomes even when those outcomes are beyond their control, Hypothesis 1 suggests that people will be more curious when the stakes are higher.

Hypothesis 2 holds because curiosity tends to be an increasing function of the *salience* of the information gap. For example, consider an employee's desire to know the results of her annual performance evaluation. If the evaluation was conducted a long time ago, it is no longer very salient, and she may have already forgotten about it. However, if the evaluation took place recently (and especially if the employee's supervisor already knows the results and has scheduled a meeting to discuss them), it would be much more salient. We would thus predict that the employee would be more curious about her performance feedback in the latter scenarios. Note that although we generally believe that time delay decreases salience, the length of the delay may be relevant. In conversation or advertising, salience may be heightened by a well-timed "pregnant" pause. Indeed, people report greater curiosity after such pauses (Kupor and Tormala 2015).

Hypothesis 3 holds because motivated attention generates stronger desire for information as the valence of anticipated beliefs increases. Returning to the example of the performance feedback, if the employee is uncertain whether she got the *best* evaluation among her colleagues or "merely" an excellent rating (i.e., all possible answers have positive valence), then she would enjoy looking at (and thinking about) her performance report. On the other hand, if she is uncertain whether she failed to meet expectations along a single criterion or disappointed her supervisor in multiple ways (i.e., all possible answers have negative valence), then looking at (and thinking about) her report is likely to be unpleasant. The employee of the month is likely to eagerly review her report (over and over again), whereas the struggling worker is more likely to quickly dispose of it without a glance or hide it in a place where it will hopefully be forgotten. Hypothesis 3 makes sense of a variety of existing empirical findings. For example, willingness to pay for an assessment of one's intelligence or beauty (relative to others) increases as one's subjective prior belief about this assessment becomes more favorable (Eil and Rao 2011, Burks et al. 2013, Möbius et al. 2022). The well-documented "ostrich effect" is the finding that investors 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 (Karlsson et al. 2009, Sicherman et al. 2015). Similarly, people are more likely to look up the value of their bank accounts immediately after getting paid (Olafsson and Pagel 2017).

3. Experimental Evidence

We report data from five online experiments aimed at testing the three predictions presented in Section 2 (see Table 1). In Section 3.1, we present data from two experiments that study demand for epistemic information using a paradigm in which anticipated beliefs have neutral valence, and therefore, there is no desire to avoid information. These experiments show that individuals are willing to exert effort to acquire noninstrumental information (the answer to a rebus puzzle). In this domain, we test Hypotheses 1 and 2 and find that increasing the importance or the salience of an unanswered question increases demand for finding out the answer. In Section 3.2, we again test these hypotheses using ego-relevant information (score on a test), whereby individuals could hold negative beliefs that could motivate information avoidance. On top of testing Hypotheses 1 and 2, this paradigm also lets us test Hypothesis 3—that manipulating the valence of anticipated beliefs affects demand for information.

In all experiments, we capture demand for information by measuring individuals' willingness to spend time and exert effort to obtain information. We conducted all of our experiments on an online labor market platform, Amazon Mechanical Turk, where workers can browse different tasks and choose which ones to complete in exchange for monetary compensation. In this setting, the choice of spending extra time and exerting effort merely to obtain information (without being compensated for the extra effort) is consequential, as it generates a clear opportunity cost; workers can choose to spend this extra time on other available tasks and can get compensated for them. We choose this outcome measure rather than eliciting participants' willingness to pay for information because requiring an expenditure of time and effort to obtain information creates a decision that is more naturalistic and similar to decisions encountered commonly in everyday life.

3.1. Epistemic (Neutrally Valenced) Information: The Rebus Puzzle Paradigm

We first report data from two experiments that investigate demand for information while keeping the valence of this information neutral, thereby eliminating the motive to direct attention away from bad news and allowing us to isolate the curiosity motive. We study how the importance and salience of an information gap affect the demand for information in this domain.

We designed a novel experimental paradigm in which participants could exert effort to learn the answer to a challenging rebus puzzle. In the experiments, participants first try solving two practice puzzles and then can earn a $2 bonus payment by successfully solving three bonus puzzles. The final puzzle is quite challenging (only 20% of participants could solve it in a pretest), so the majority of participants fail to solve it (and thus, fail to obtain the bonus). After completing the task, participants have the opportunity to exert effort to reveal the answer to the last puzzle. We ask participants if they are interested in learning the solution to this puzzle ("YES" or "NO"). Participants who click the "YES" button then get a pop-up message that instructs them to click the button again if they want to see the solution. They have to click "YES" a number of times (10 times in Experiment 1A; 5 times in Experiment 1B) to actually see the solution, but they do not know this number ahead of time. During each iteration, participants can choose to click "YES" to indicate they want to see the solution, or participants can choose to click "NO" to skip revealing the solution and conclude the survey immediately. Note that in each iteration, we simply instruct participants to click again if they want to see the solution, without promising them that the solution would be immediately displayed. If participants revise their expectation about the necessary number of clicks upward every time they are asked to click another time, they may experience an increasing cost associated with revealing the solution. Because participants did not know whether the repeated pop-up message was an error or a feature of the experiment, it is possible that they became increasingly frustrated after repeatedly failing to reveal the solution. We treat frustration as a potential type of cost associated with revealing the solution (in addition to effort and time), and our predictions are robust whether participants became increasingly frustrated or not. We expect participants to stop clicking when the cost exceeds the expected gain in utility from satisfying curiosity. We estimate curiosity about the solution by measuring the number

Table 1. The Experiments

Experiment	Information (valence)	Paradigm	Hypothesis	Treatments	Sample size, N
1A	Epistemic (neutral)	Rebus puzzle	Hypothesis 1: Importance	*High Importance, Low Importance*	838
1B	Epistemic (neutral)	Rebus puzzle	Hypothesis 2: Salience	*Immediate, Delayed*	157
2A	Ego relevant (valenced)	FER test	Hypothesis 1: Importance	*High Bonus, Low Bonus*	470
2B	Ego relevant (valenced)	FER test	Hypothesis 2: Salience	*Immediate, Delayed*	398
2C	Ego relevant (valenced)	FER test	Hypothesis 3: Valence	*Easy, Hard*	498

of clicks on the reveal button. A higher number of clicks reveals willingness to pay a higher nonmonetary cost and thus, implies stronger curiosity.

3.1.1. The Experiments.

3.1.1.1. Experiment 1A: Importance. In the first experiment, we test whether increasing the perceived importance of an information gap affects demand to learn the solution to the puzzle. We manipulate attention toward the last puzzle by varying whether knowing the answer to this puzzle is important for obtaining the bonus. To do so, we vary the order in which the five puzzles are presented, effectively varying which puzzles are for practice and which count for payment (see Figure 1). The last (especially challenging) puzzle in the sequence ("moral support") is the same across treatments.

The two treatments thus attempt to vary whether knowing the solution to this last puzzle is pivotal for getting the bonus. In the *High Importance* treatment, two of the three bonus puzzles are easy to solve (93% and 90% managed to solve the "jack-in-the-box" and the "falling asleep" puzzles in a pretest, respectively), making the last puzzle usually pivotal (i.e., important for getting the $2 bonus). In the *Low Importance* treatment, the last puzzle follows two other challenging bonus puzzles (only 37% and 30% could solve the "green fingers" and the "painless operation" puzzles in a pretest, respectively), typically making none of them pivotal (unilaterally important). The two practice puzzles in each treatment are the first two bonus puzzles from the other treatment to ensure that participants in both treatments complete the same set of five puzzles, keeping the required effort and the overall difficulty constant.

We hypothesize that, conditional on failing to solve the last puzzle (and thus, failing to get the bonus), the answer to the last puzzle is perceived as more important if it is the only answer that prevented one from winning the bonus (*High Importance* treatment) as opposed to being one of several answers that prevented individuals from getting the bonus (*Low Importance* treatment). We then predict that making this puzzle more important attracts more attention to the information gap regarding the correct answer to the puzzle, increasing the desire to find out the solution.

According to the model, the answer to a pivotal puzzle is important because it has the chance to affect the final payoff, which would affect utility. The possibility of getting the bonus induces spread in the potential utility depending on the answer. However, in our design, participants learn that they failed to receive the bonus right before being asked whether they want to find out the solution to the puzzle. Hence, finding out the solution cannot affect their payoff in any way. Yet, the assumption that importance does not change until a person adapts to a new belief implies that participants may still consider the answer

Figure 1. (Color online) Experimental Procedure, Experiment 1A

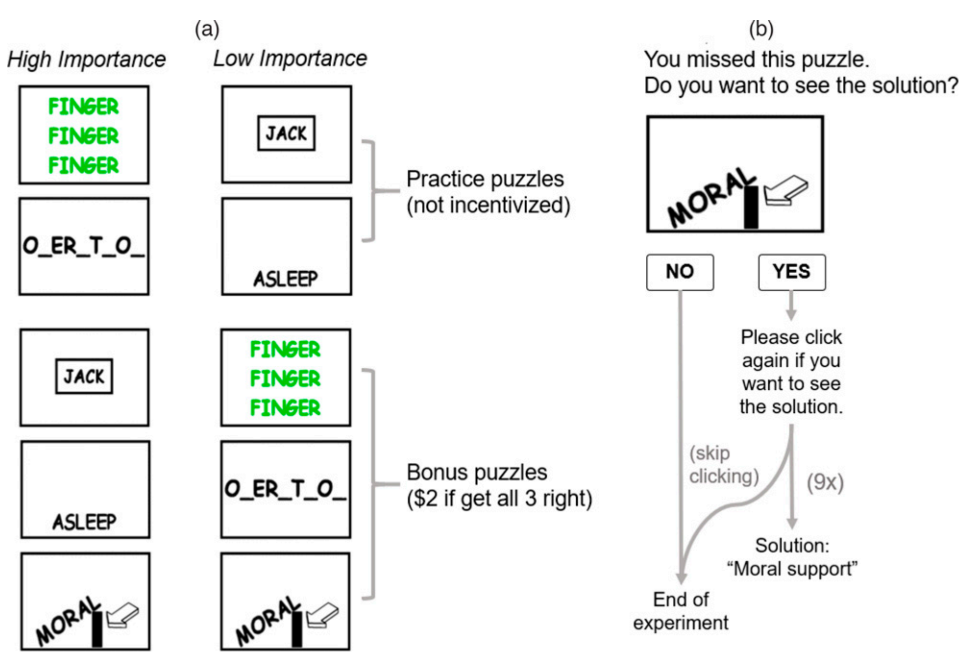

Notes. Panel (a) shows the two experimental treatments of Experiment 1A: *High Importance* (left panel) and *Low Importance* (right panel). Panel (b) shows the behavioral measure used in Experiment 1A: the willingness to click to reveal the solution. If participants clicked "YES" initially and clicked nine more times, the solution was revealed. If at any point they clicked "NO," the experiment concluded without revealing the solution. This measure was the same across treatments.

important, even while recognizing that it no longer has instrumental value. So, we predict higher curiosity in the *High Importance* treatment than in the *Low Importance* treatment.

3.1.1.2. Experiment 1B: Salience. In the second experiment, we test the prediction that curiosity is stronger for information gaps that opened more recently and thus, are more salient. In the experiment, participants complete two tasks on two consecutive days. On the first day, they work on the rebus puzzle task as in Experiment 1A. They try two practice puzzles and can then earn a $2 bonus payment for correctly solving all three of the subsequent bonus puzzles. All participants face the same sequence of puzzles as in the *High Importance* treatment of Experiment 1A. On the second day, participants answer a set of knowledge trivia questions, which serves as a filler task, unrelated to the puzzle task. As in Experiment 1A, participants have the option to exert effort to reveal the solution to the "moral support" puzzle, and curiosity is elicited the same way. To vary salience, we manipulate the timing of when participants are asked if they would like to reveal the solution. In the *Immediate* treatment, participants are asked immediately after completing the puzzle task on the first day whether they would like to reveal the solution immediately. In the *Delayed* treatment, participants are asked only on the second day (i.e., a day after completing the puzzle task) whether they would like to reveal the solution at that time. We predict that, conditional on failing to solve the "moral support" puzzle, participants will be more curious to learn the answer when this information gap is more salient (i.e., opened just before the opportunity to obtain the information, as opposed to a day earlier).[8]

3.1.1.3. Procedures. We recruited participants from Amazon Mechanical Turk. Participants received a fixed payment of $0.25 for a three-minute study and had a chance to earn a $2 bonus if they solved all three bonus puzzles. In the Salience experiment (1B), we promised participants an additional $1.50 for completing the second stage of the experiment a day after completing the first. The instructions for both experiments are available in Online Appendix B.

In both experiments, participants knew that their performance in the practice rounds did not affect their final payment. Upon enrollment, we randomly assigned participants in the Importance experiment (1A) to the *High Importance* or the *Low Importance* treatment and participants in the Salience experiment (1B) to the *Immediate* or *Delayed* treatment. After completing the two practice puzzles and the three bonus puzzles, participants received feedback about their performance (how many of the practice and bonus puzzles they solved correctly) and whether they won the bonus. In the Importance experiment and in the *Immediate* treatment of the Salience experiment, we then displayed the last (moral support) puzzle again and told participants whether they successfully solved it. Next, we elicited their curiosity by asking participants if they wanted to see the solution ("YES" or "NO") and repeatedly asking this question if they clicked "YES" as described. Participants who clicked the "NO" button ended the study without revealing the solution. Because the puzzle comes in the form of an image, it was hard for participants to find the solution elsewhere (e.g., on the internet). In the Salience experiment, we invited participants in both treatments to the second stage via email the day after the first stage. In the second stage, participants worked on a set of trivia questions followed by some demographics questions. Their compensation did not depend on performance—they all received a fixed payment of $1.50. Whereas in the *Immediate* treatment, participants were prompted to find out the solution of their last rebus puzzle ("moral support") at the end of stage 1, in the *Delayed* treatment we prompted participants to find out the solution one day later, at the end of stage 2.

3.1.2. Results.

3.1.2.1. Experiment 1A: Importance. A total of 853 participants (90.4%) completed the experiment. We excluded 15 participants (1.8%) who submitted duplicate responses. The final sample contained 838 participants (45.1% female): 418 in the *High Importance* treatment and 420 in the *Low Importance* treatment. We determined these sample sizes by conducting an a priori power analysis (see Online Appendix C).[9]

3.1.2.1.1. Performance on the Puzzle Task. There was no significant difference in the proportion of people who failed to solve the "moral support" puzzle: 236 people (56.5%) missed this puzzle in the *High Importance* treatment, and 239 (56.9%) missed it in the *Low Importance* treatment, $\chi^2(1, N = 838) = 0.004$, $p = 0.952$. Participants who failed to answer the "moral support" puzzle were able to solve significantly more bonus puzzles in the *High Importance* treatment, in which two easy puzzles preceded the last puzzle, $M = 1.66$, than in the *Low Importance* treatment, in which two hard puzzles preceded the last puzzle, $M = 0.49$, $t(469) = 19.619$, $p < 0.001$, Cohen's $d = 1.80$, 95% confidence interval (CI) [1.05, 1.28]. In the *High Importance* treatment, 183 people (77.5%) solved both bonus puzzles before the last puzzle; thus, they missed their $2 bonus only because of missing the last puzzle. In the *Low Importance* treatment, 223 (93.3%) failed to solve at least one of the other two bonus puzzles before the last puzzle, missing the bonus because of missing multiple puzzles. In terms of overall performance, participants' total scores (practice + bonus) were not significantly different between the *High Importance*, $M = 2.03$, and the *Low*

Importance treatments, $M = 2.08$, $t(467) = 0.532, p = 0.595$, Cohen's $d = 0.05$, 95% CI $[-0.23, 0.13]$. Similarly, there was no difference in total time spent trying to solve the puzzles (practice + bonus): $M = 2.82$ minutes in the *High Importance* treatment and $M = 3.05$ minutes in the *Low Importance* treatment, $t(466) = 1.085, p = 0.278$. Thus, the manipulation does not appear to have affected participants' effort, which suggests that treatment effects were driven by attention upon realizing the "moral support" puzzle was pivotal, as opposed to the effort put into solving it.

3.1.2.1.2. Main Results: Willingness to Exert Effort to Reveal the Solution. Because the information gap of our interest existed only for those participants who could not answer the "moral support" puzzle—and not for those who managed to solve it—our key analyses focus on the group of people who failed to answer this puzzle correctly.[10] Of the 475 participants who could not solve the "moral support" puzzle, only 32 (6.7%) declined to reveal the solution immediately—that is, to click on "YES" even a single time—and the average number of clicks was $M = 5.00$ (standard deviation (SD) = 3.10). We report the detailed distribution of click counts by experimental treatment in Figure D1 in Online Appendix D.

Crucially, participants in the *High Importance* treatment clicked significantly more, $M = 5.68$, than in the *Low Importance* treatment, $M = 4.34$, $t(472) = 4.814$, $p < 0.001$, Cohen's $d = 0.44$, 95% CI $[0.79, 1.89]$ (see Figure 2(a)). Participants in the *High Importance* treatment were significantly more likely to start clicking to find out the solution, $M = 97.9\%$, compared with participants in the *Low Importance* treatment, $M = 88.7\%$ ($\chi^2(1, N = 475) = 14.494, p < 0.001$, Cohen's $w = 0.183$) and to click the 10 times to actually reveal the solution ($M = 25.4\%$ for *High Importance*, $M = 14.2\%$ for *Low Importance*, $\chi^2(1, N = 475) = 8.688, p = 0.003$, Cohen's $w = 0.141$).

We also investigate these results using ordinary least squares (OLS) regression analyses. Table D1 in Online Appendix D shows that the overall number of clicks (column (1)), the proportion of participants who clicked at all (column (3)), and the proportion of participants who clicked 10 times (column (5)) were all significantly higher in the *High Importance* treatment. These results are robust to controlling for participants' total score, the time they spent on each puzzle, and their gender (columns (2), (4), and (6)). The total score and the time spent on the "moral support" puzzle also significantly predict the willingness to reveal the solution. These results are also consistent with our prediction regarding importance. The more invested people were—that is, the more effort they put into solving the puzzles—and the closer they got to solving all the puzzles, the more important the solution to the last puzzle was.

3.1.2.2. Experiment 1B: Salience. A total of 200 participants (90.9%) completed the first stage of the experiment and were invited the next day to participate in the second stage. Of these, 74.5% ($n = 164$) completed the second stage. Among those who completed both stages of the experiment, we excluded seven duplicate responses (4.3%), leaving us with $n = 157$ observations (49.0% female). The proportion of participants who completed the second stage does not differ by treatments ($n = 77$ participants in the *Immediate* treatment and $n = 80$ participants in the *Delayed* treatment), $\chi^2(1, N = 200) = 0.119$, $p = 0.731$.

3.1.2.2.1. Performance on the Puzzle Task. There was no significant difference in the performance of participants between the two treatments. People solved on average 2.04 and 2.06 puzzles in the *Immediate* and *Delayed* treatments, respectively, $t(154) = 0.183, p = 0.855$, Cohen's $d = 0.03$, 95% CI $[-0.28, 0.23]$. Across treatments, 108 people (68.8%) failed to solve the "moral support" puzzle. There was no significant difference in the proportion of people who failed to answer this puzzle correctly between the two treatments; 55 people (71.4%) missed this puzzle in the *Immediate* treatment, and 53 (66.3%) missed it in the *Delayed* treatment, $\chi^2(1, N = 157) = 0.279, p = 0.598$.

3.1.2.2.2. Willingness to Exert Effort to Reveal the Solution. Of the 108 participants who could not solve the "moral support" puzzle, only 8 (7.4%) declined to click on "YES" even a single time, and the average number of clicks was $M = 3.76$ ($SD = 1.53$). We report the detailed distribution of click counts by experimental treatment in Figure D2 in Online Appendix D.

Key to our hypothesis, participants in the *Immediate* treatment clicked significantly more times, $M = 4.27$, than in the *Delayed* treatment, $M = 3.23$, $t(89) = 3.727$, $p < 0.001$, Cohen's $d = 0.72$, 95% CI $[0.49, 1.60]$ (see Figure 2(b)).

Participants in the *Immediate* treatment were also significantly more likely to start clicking to find out the solution, $M = 98.2\%$, compared with participants in the *Delayed* treatment, $M = 86.8\%$, $p = 0.030$.[11] Finally, participants in the *Immediate* treatment revealed the solution (i.e., clicked five times) significantly more often, $M = 63.6\%$, compared with participants in the *Delayed* treatment, $M = 37.7\%$, $\chi^2(1, N = 108) = 6.246$, $p = 0.012$, Cohen's $w = 0.259$.

In Table D2 in Online Appendix D, we confirm these results using OLS regression, where we control for other factors that could potentially explain differences in the willingness to reveal the solution, such as the total score achieved, the time spent on the puzzles, and demographic covariates. The main results are robust to the inclusion of additional predictors. As in Experiment 1A, the total score and the time spent on solving puzzles also predict the willingness to reveal the solution. That is, the more effort people put into

Figure 2. (Color online) Main Results, Experiments 1A and 1B

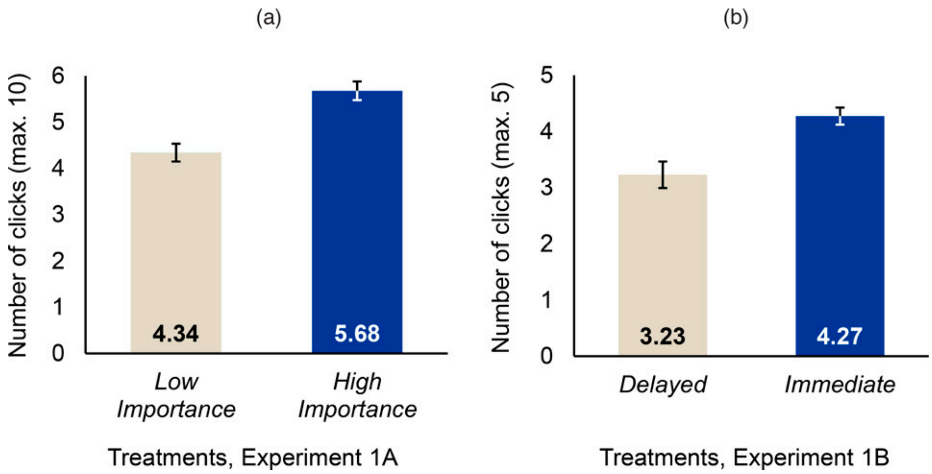

Notes. (a) Experiment 1A. (b) Experiment 1B. Error bars represent ±1 standard error.

solving the puzzles, the more curious they were about the solution to the last puzzle.

Because our analyses were limited to the sample of participants who completed *both* stages of the experiment—as opposed to comparing everyone in the *Immediate* treatment who completed the first stage with the group of participants in the *Delayed* treatment who completed both stages—the difference in the willingness to reveal the solution between treatments cannot be explained by a selection effect. We also note that we do not find any evidence of a selection effect (see Online Figure D3 and additional analyses in Online Appendix D).

3.1.3. Discussion. Experiments 1A and 1B provide support for the hypotheses that increasing the perceived importance and salience of an information gap increases demand for information (Hypotheses 1 and 2, respectively). The puzzle task allows us to isolate the curiosity motive from the motive to avoid negative information. However, these experiments rely on some nontrivial assumptions.

First, we assume that participants would not be able to find out the solution to the puzzle on their own by searching online. As discussed, the puzzles are pictures, which made searching for their solutions difficult. However, if participants did somehow find the solution on their own in the *Delayed* treatment in Experiment 1B, this would have reduced their willingness to reveal the solution the next day, apart from any salience effect. Note, however, that the decrease in clicking in the *Delayed* treatment extends beyond the increased proportion who immediately decline to reveal the solution (see Figure D2 in Online Appendix D).

Second, in Experiment 1A we study information acquisition after the information loses its instrumental value, relying on the assumption that finding out the answer would not immediately cease to be important after participants found out that they did not win the prize. This is likely a valid assumption because such hedonic adaptation is usually not instantaneous (Wilson et al. 2005), and we expected participants to continue to be curious about answers that were clearly important a moment ago.

To allay potential concerns that our findings rely on the validity of these assumptions, in Section 3.2 we report data from additional experiments in which we study demand for information using a different paradigm. In these experiments, participants can exert effort to learn ego-relevant information. This different paradigm allows us to manipulate importance without any additional assumption about adaptation and to provide participants with information that they cannot obtain anywhere else. With this paradigm, we provide additional supportive evidence for Hypothesis 1 (Experiment 2A) and Hypothesis 2 (Experiment 2B), showing the robustness of these effects in a domain where anticipated beliefs can be negative. This paradigm also allows us to test how valence affects information acquisition (Experiment 2C).

3.2. Ego-Relevant (Valenced) Information: The Facial Expression Recognition Test Paradigm

We now report data from three experiments that investigate demand for ego-relevant information using a paradigm in which individuals can learn about their own performance on a test. In this domain, the curiosity motive may sometimes be overwhelmed by the desire to direct attention away from bad news. We first show that, even in this domain, directing attention to the information gap by making it feel more important or salient increases demand for information (Experiments 2A and 2B, respectively). Furthermore, we then show that the perceived valence of the information gap affects demand for information; individuals are more likely to acquire information when valence is more positive (Experiment 2C).

In these experiments, participants complete an online facial expression recognition (FER) test we designed, which measures individuals' ability to recognize emotions from facial expressions. The FER test presents individuals with a sequence of 40 photos of faces and asks them to guess which of six emotions (happiness, sadness, anger, disgust, fear, surprise) the people in the photos are displaying (see Figure 3 and Online Appendix B for the stimuli and detailed instructions).

By taking the FER test, participants activate an information gap about their ability to recognize emotions. Information about their performance may be ego relevant if they want to believe they are good at recognizing emotions. We chose this test because we believe that participants have little prior knowledge about their ability to recognize emotions and would naturally be curious about their performance.

In order to reveal their scores and thus, to close the information gap about their ability to identify emotions, participants have to complete an additional task for no extra payment. Specifically, after completing the FER test, participants choose whether to complete a boring three-minute extra task incentivized only by the promise that if they do it, they can find out their *exact* score on the FER test and their relative ranking (percentile) compared with other participants. As we designed a unique FER test specifically for these experiments, participants could not find out their score on it elsewhere, only by completing the three-minute extra task we offered.[12] We estimate demand for information by measuring the fraction of individuals who are willing to start and complete this extra task in order to learn their score.

3.2.1. The Experiments.

3.2.1.1. Experiment 2A: Importance. In this experiment, we investigate how a manipulation of importance affects willingness to exert effort to learn about one's score on the FER test. Differently from Experiment 1A, we manipulate importance here by directly altering the size of a bonus participants can receive for doing well on the FER test. Prior to completing the task, participants learn that they will receive a bonus for correctly solving 50% of the tasks (20 of 40 photos). A larger potential bonus makes their score more important because it makes participants care more about whether they get the bonus. That is, it induces greater spread in their potential utility depending on their score, as they would be happier if their score gets them the bonus and more disappointed if not. After finishing the test, they are told that they will find out whether they got the bonus, and they are also given the option to work on the extra three-minute task to learn their exact score. They know they will learn whether they got the bonus, regardless of their choice to learn their exact score, and getting this additional information cannot change whether they will get the bonus. Thus, we can identify demand for information about their exact score over and above their desire to know if they earned the bonus.

In the *High Bonus* treatment, the bonus is \$1. This bonus is on top of a \$0.75 fixed payment for completing the task. In the *Low Bonus* treatment, the bonus is \$0.05. We randomly assign *Low Bonus* participants to receive either the same fixed payment as participants in the *High Bonus* treatment (\$0.75) or a higher fixed payment of \$1.70. Conditional on scoring above threshold, participants in this *Low Bonus (high fixed pay)* treatment receive the same total earnings (\$1.70 + \$0.05 = \$1.75) as participants in the *High*

Figure 3. (Color online) The Facial Expression Recognition Test, Experiments 2A–2C

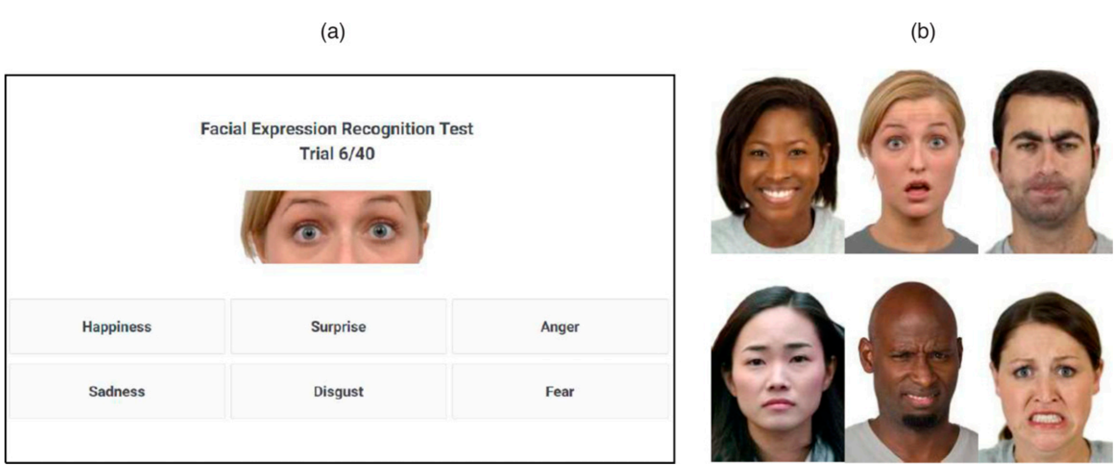

Notes. (a) Participants were shown cropped photos and had to guess which of the six emotions the person in the photo was experiencing. (b) A sample selection of six uncropped photos, representing the six emotions (happiness, sadness, anger, disgust, fear, and surprise). We report the full list of photos in Online Appendix B.

Bonus treatment ($0.75 + $1 = $1.75). We predict that the proportion of individuals who choose to start and complete the extra task to learn their exact score will be higher in the *High Bonus* than in the *Low Bonus* treatment.

To further support our theoretical predictions and to highlight the mechanism that drives information preferences, we include three manipulation check questions that we ask after the extra task but before participants learn their outcome (and exact score if they decided to complete the extra task). We ask how happy/unhappy (on a scale from –100 to + 100) participants would feel (a) if they got the bonus (U_{win}) and (b) if they did not get the bonus (U_{lose}). Further, we elicit beliefs about the likelihood of getting the bonus (p_{win}: 0%–100%). These three measures allow us to calculate the *standard deviation* of the anticipated happiness (i.e., utility):

$$\mathrm{SD}(U) = \sqrt{(p_{win})(U_{win} - \bar{U})^2 + (1 - p_{win})(U_{lose} - \bar{U})^2},$$

where $\bar{U} = (p_{win})U_{win} + (1 - p_{win})U_{lose}$. According to our theory, the more spread out the distribution of possible utilities is (e.g., the higher the standard deviation of anticipated happiness), the more important the information gap is. If SD(U) is higher in the *High Bonus* treatment than in the *Low Bonus* treatments, then participants feel that their score is more important when they can win a larger bonus (i.e., the experimental manipulation is successful).[13]

3.2.1.2. Experiment 2B: Salience. In this experiment, we use the FER paradigm to test the hypothesis that individuals are more curious about information gaps that were opened more recently (i.e., that are more salient). The experiment is similar to Experiment 2A, with the following exceptions. (1) Participants' performance on the test is not incentivized. (2) Participants do not receive any immediate feedback on their performance. Instead, they receive a follow-up email with an opportunity to complete an extra three-minute task to learn about their performance on the FER test. The extra task does not result in any payment. As in Experiment 2A, participants are not forced to complete the extra task if they start it. However, if they quit, they do not receive *any* feedback about their score on the test. We manipulate salience by varying how recently the information gap is opened before the opportunity to complete the extra task. In the *Immediate* treatment, participants receive the email within 15 minutes after the FER test. In the *Delayed* treatment, participants receive it 24 hours after they complete the test. We predict that the proportion of individuals who choose to complete the test will be higher in the *Immediate* than in the *Delayed* treatment.

3.2.1.3. Experiment 2C: Valence. In this experiment, we use the FER test to investigate the prediction that people are willing to exert more effort to fill an information gap when the possible answers have more positive valence. To manipulate valence, we manipulate the difficulty of the task, thereby affecting participants' performance. In the *Easy* treatment, we oversample easy photos. In the *Hard* treatment, we oversample difficult photos. After performing the task, participants receive preliminary information about their performance. Specifically, we show participants the distribution of scores on a prior FER test of moderate difficulty, as well as the average score, and highlight two potential scores they could have gotten in the test, informing them that one of the two is (truthfully) their actual score. We manipulate whether the alternative score is 20 points higher or 20 points lower than the actual score. Depending on our treatments and their performance, three scenarios are possible: (1) *Good Expected News*: both scores are better than average (i.e., the participant could receive only good news by revealing which is the true score) (Figure 4, right panel); (2) *Bad Expected News*: both scores are worse than average (i.e., only bad news) (Figure 4, left panel); or (3) *Mixed Expected News*: one score is at least average or better than that, whereas the other is at most average or worse than that (i.e., mixed news: either good or bad). Although people may be naturally curious about their score in all three scenarios, varying the valence affects the motive to direct attention away from negative news (and toward positive news). After receiving the preliminary information, participants have the option to complete the three-minute real effort task to find out which of the two highlighted scores is their actual score.

3.2.1.4. Procedures. We recruited participants on Amazon Mechanical Turk. In all experiments, participants received a fixed payment of $0.75 for completing the task, with the exception of half of the participants in the *Low Bonus* treatment who received a $1.70 fixed payment instead. For each experiment, we determined the sample sizes by conducting an a priori power analysis; we report all power analyses in Online Appendix C.

In the Importance experiment (2A), which was preregistered on aspredicted.org, on top of the fixed payment, participants received an additional bonus for successfully completing 50% of the task: $0.05 in the *Low Bonus* treatment and $0.75 in the *High Bonus* treatment.

After completing the test but before learning whether they received the bonus, participants had the opportunity to reveal their exact score and their relative ranking (percentile) compared with other participants by completing a boring three-minute extra task. In

Figure 4. (Color online) Experimental Stimuli, Experiment 2C

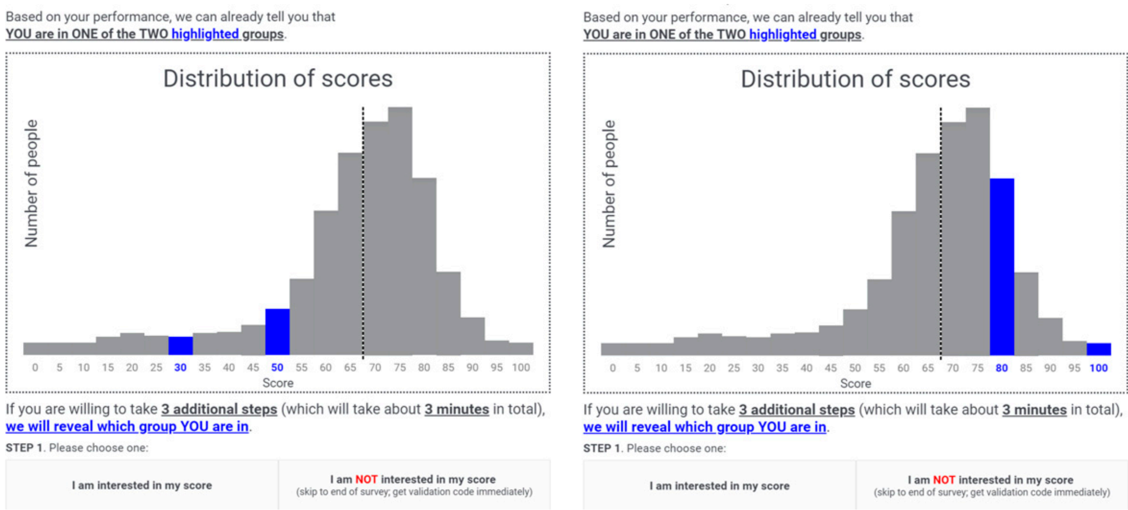

Notes. The sample screens depict the page where participants were told that their score is in one of the two highlighted bins and were offered the opportunity to reveal their scores. (Left panel) Both the actual score and the alternative score are worse than average (bad expected news). (Right panel) Both the actual score and the alternative score are better than average (good expected news).

this extra task, they had to guess the age of 15 people and indicate their confidence in their guesses. We included an attention check question at the end of the FER test. After the three manipulation check questions on anticipated happiness and probability of getting the bonus, we also included a comprehension check question testing whether participants understood that revealing (or not revealing) their score would not affect their payment. As preregistered, we excluded participants who failed either the attention check or the comprehension check.

In the Salience experiment (2B), participants completed the FER test and did not receive any performance-based bonus. After completing the test, they received a follow-up email informing them of the opportunity to reveal their score. In the *Immediate* treatment, the email was sent within 15 minutes of completing the FER test. In the *Delayed* treatment, the email was sent 24 hours after participants completed the test. The email said: "Thank you for taking the Facial Expression Recognition Test on [DATE]. Now you have the opportunity to learn your Facial Expression Recognition Score! If you are willing to take 3 additional steps (which will take about 3 minutes in total), we will reveal your FER Score, and you will also see how well you did compared to other people. To reveal your score, please open the link and follow the instructions: [LINK]." In order to control for any time of the day/day of the week effects, we send these emails to everyone *at the same time*, but participants in different treatments *completed the FER test at different times*. After opening the link in the email, participants could complete the same three-minute extra task as in Experiment 2A in order to learn their score on the FER test and their relative ranking (percentile) compared with other participants. There were no attention checks or comprehension questions in this experiment.

In the Valence experiment (2C), participants completed either an easy or hard version of the FER test (*Easy* and *Hard* treatments, respectively). They were then shown the distribution of scores on an earlier (moderate difficulty) FER test, showing the proportion of people in each of the 21 score bins (0%–100% in 5% increments) and also indicating the average score (see Figure B1 in Online Appendix B). Then, we highlighted two of these score bins, one of which contained their actual score and another bin that contained an alternative score either 20 points higher or 20 points lower than their actual score.[14] Participants could then complete the same three-minute extra task as in Experiments 2A and 2B to reveal which one was their score.

In all experiments, we measured the proportion of people who started and completed this task, and we also recorded participants' gender and age.

3.2.2 Results.

3.2.2.1. Experiment 2A: Importance. A total of 632 participants (95.6%) completed the experiment. As preregistered, we excluded 93 participants (14.7%) who failed the comprehension check question, 68 participants (10.8%) who failed the attention check question, and 1 participant (0.2%) who submitted a duplicate response. The final sample contained 470 participants (47.4% female, $M_{age} = 41.7$ years): 163 in the *High Bonus* treatment, 154 in the *Low Bonus (regular fixed pay)* treatment, and 153 in the *Low Bonus (high fixed pay)* treatment.

3.2.2.1.1. Manipulation Checks and Subjective Importance. Consistent with our intended manipulation, participants reported higher expected happiness in relation to getting the bonus in the *High Bonus* treatment, $M = 76.2$, than in both the *Low Bonus (regular fixed pay)* treatment, $M = 44.5$, $t(283) = 8.781, p < 0.001$, Cohen's $d = 0.99$, 95% CI [24.62, 38.84], and in the *Low Bonus (high fixed pay)* treatment, $M = 39.5$, $t(287) = 10.379$, $p < 0.001$, Cohen's $d = 1.18$, 95% CI [29.75, 43.68]. The expected happiness ratings upon getting the bonus were similar in the two *Low Bonus* treatments, $t(305) = 1.227$, $p = 0.221$, Cohen's $d = 0.14$, 95% CI [−3.01, 12.98].

Similarly, participants reported that they would be significantly *less* happy if they *would not* get the bonus in the *High Bonus* treatment, $M = -42.8$, than in both the *Low Bonus (regular fixed pay)* treatment, $M = -13.6$, $t(315) = 7.143, p < 0.001$, Cohen's $d = 0.80$, 95% CI [−37.30, −21.19], and in the *Low Bonus (high fixed pay)* treatment, $M = -16.5$, $t(305) = 6.857, p < 0.001$, Cohen's $d = 0.77$, 95% CI [−33.95, −18.81]. Again, there was no significant difference in the happiness ratings between the two *Low Bonus* treatments, $t(299) = 0.774$, $p = 0.440$, Cohen's $d = 0.09$, 95% CI [−4.42, 10.16].

Importantly, the experimental manipulation only affected participants' expectation about their (un-)happiness upon getting (or not getting) the bonus but not their expectation about the *likelihood* of winning. Participants reported that they would be equally likely to win in the *High Bonus* treatment, $M = 63.5\%$, as in the *Low Bonus (regular fixed pay)* treatment, $M = 64.4\%$, and in the *Low Bonus (high fixed pay)* treatment, $M = 65.6\%$. We observed no significant differences between any treatments, all $p > 0.341$. Thus, participants did not expect to receive better or more informative news in the *High Bonus* treatment than in the *Low Bonus* treatments, ruling out alternative explanations based on valence or expected informativeness.

Finally, we calculated SD(U) for each participant from the three measures reported (p_{win}, U_{win}, and U_{lose}). Consistent with our intended manipulation, the subjective importance, as measured by SD(U), was significantly higher in the *High Bonus* treatment, $M = 49.0$, than in the *Low Bonus (regular fixed pay)* treatment, $M = 23.9$, $t(312) = 9.126, p < 0.001$, Cohen's $d = 1.02$, 95% CI [19.70, 30.53], and in the *Low Bonus (high fixed pay)* treatment, $M = 21.7$, $t(312) = 9.78$, $p < 0.001$, Cohen's $d = 1.10$, 95% CI [21.81, 32.79]. There was no significant difference between the two *Low Bonus* treatments, $t(305) = 0.841, p = 0.401$, Cohen's $d = 0.10$, 95% CI [−2.93, 7.31].

These manipulation checks confirm that the information about one's performance on the FER test in the *High Bonus* treatment was deemed to be more important than the same information in the *Low Bonus* treatments. In addition, we did not observe any significant differences between the two *Low Bonus* treatments, which indicates that participants treated the information as about equally important in both of these treatments. Therefore, we decided to pool these two treatments in subsequent analyses.[15]

3.2.2.1.2. Main Results: Exerting Effort to See the Solution. A significantly higher proportion of participants started the extra task in the *High Bonus* treatment, $M = 46.6\%$, than in the *Low Bonus* treatments, $M = 33.6\%$ ($M_{regular} = 38.3\%$, $M_{high} = 28.8\%$), $\chi^2(1, N = 470) = 7.17, p = 0.007$, Cohen's $w = 0.128$ (see Figure 5(a)). Similarly, a significantly higher proportion of participants completed the extra task (and revealed their score) in the *High Bonus* treatment, $M = 42.3\%$, than in the *Low Bonus* treatments, $M = 30.9\%$ ($M_{regular} = 34.4\%$, $M_{high} = 27.5\%$), $\chi^2(1, N = 470) = 5.59, p = 0.018$, Cohen's $w = 0.114$.

To test whether the derived measure of subjective importance SD(U) predicts participants' decision to start and complete the extra task, as well as to control for potential wealth effects, actual performance, and demographic factors, we conducted hierarchical OLS regression analyses. In these models, we included the proportion of people starting and completing the extra task as dependent measures and added experimental treatment, subjective importance, level of fixed payment, actual score, age, and gender as potential predictors and covariates (see Table D3 in Online Appendix D).

This regression analysis revealed that the derived measure of subjective importance significantly predicts both whether someone starts and completes the extra task: $\beta = 0.003$, standard error (SE) = 0.001, $t(467) = 3.377$, $p < 0.001$ and $\beta = 0.003$, SE = 0.001, $t(467) = 3.232$, $p = 0.001$ for starting and completing, respectively. Furthermore, this measure not only significantly predicts participants' choice to reveal their score, but adding this measure to the model makes the experimental treatment variable become nonsignificant, $p = 0.339$ and $p = 0.464$, for starting and completing, respectively (Online Appendix D, Table D3, columns (2) and (5)). Finally, the results are robust to the inclusion of additional controls, including the amount of fixed pay, actual performance, and demographics (Online Appendix D, Table D3, columns (3) and (6)).

3.2.2.1.3. Mediation Analyses. As a final test of the proposed mechanism of subjective importance (i.e., spread of expected happiness), we conducted a mediation analysis to assess whether the effect of experimental condition on participants' desire to start and complete the extra task was mediated by the spread of their expected happiness. We included the experimental condition as the predictor variable and the proportions starting and completing the extra task as outcome variables. We then added SD(U) as the proposed mediator. All variables were standardized before conducting the mediation analysis. A bootstrapped mediation with 5,000 replications revealed

Figure 5. (Color online) Main Results and Mediation Analysis, Experiment 2A

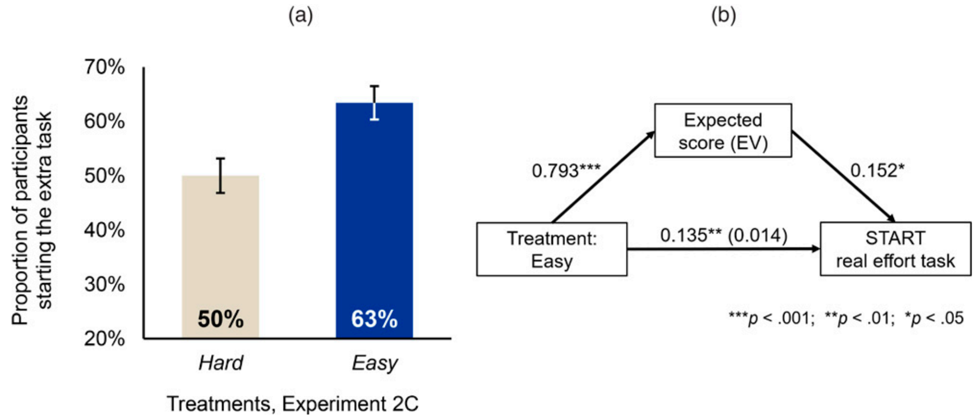

Notes. Error bars in panel (a) represent ±1 standard error. Coefficients in panel (b) are standardized β-coefficients. EV, expected value.

that subjective importance, SD(U), *fully* mediates the effect of experimental treatment on both starting and completing the extra task, $\beta = 0.080$, 95% CI [0.031, 0.132], $p < 0.001$ for starting the task (see Figure 5(b)) and $\beta = 0.076$, 95% CI [0.028, 0.130], $p < 0.001$ for completing the task.

3.2.2.2. Experiment 2B: Salience. A total of 398 participants (98.3%) completed the study (41.2% female; $M_{age} = 34.6$ years): 199 in the *Immediate* treatment and 199 in the *Delayed* treatment. We did not exclude any participants among people who completed the study.

3.2.2.2.1. Main Results. Pooling across both treatments, 165 people (41.5%) opened the link in the email to start the extra task. To test whether the experimental manipulation had an effect on the willingness to start the extra task, we first compared the proportion of people who started the task *any time* after receiving the email. Note that we collected responses to the follow-up survey for *one week* after sending the follow-up emails, allowing the participants to start the extra task any time within one week following the test. The overwhelming majority of people (95%) who ever opened the link to the extra task did so within the first eight hours after receiving the follow-up email, and no one started the follow-up survey more than two days after receiving the email.

Although 91 people (45.7%) started the extra task in the *Immediate* treatment, only 74 people (37.2%) did so in the *Delayed*, although this difference is not quite significant, $\chi^2(1, N = 398) = 2.650$, $p = 0.104$, Cohen's $w = 0.087$. Because some participants did not check their email right away, thereby experiencing significantly longer delays than the 15-minute delay we intended for them, we also repeat the analyses by comparing subsets of participants *from both treatments* who opened the study link within a specific length of time after receiving it. We expected a bigger difference between experimental treatments when examining shorter time frames—for example, people opening the study link within the first hour after receiving the email, when the information gap was still likely to be salient in the *Immediate* treatment.

Although 59 people (29.7%) started the extra task within one hour after receiving the email in the *Immediate* treatment, only 38 people (19.1%) did so in the *Delayed* treatment, $\chi^2(1, N = 398) = 5.453, p = 0.020$, Cohen's $w = 0.123$. By contrast, when comparing the proportion of people who started the extra task *at least* one hour later after receiving the email, we do not find any significant difference—the proportions are virtually identical: 16.1% in the *Immediate* treatment and 18.1% in the *Delayed* treatment, $\chi^2(1, N = 398) = 0.160, p = 0.690$.

The boundary between "shorter" and "longer" time frames is, however, somewhat subjective. We thus check the robustness of this result at different thresholds: 15 minutes, 30 minutes, two hours, and four hours. Importantly, we obtain the same results if we set a different threshold. The proportion of people starting and completing the extra task is always higher in the *Immediate* treatment than in the *Delayed* treatment (we report the results of these analyses in Online Appendix D). These effects are almost entirely driven by the difference in the behavior of people who started the extra task within the earliest time window (within 15 minutes) after receiving the follow-up email (see Figure 6), which is consistent with the model's prediction that individuals are most curious when the information gap is the most salient.

In Table D4 in Online Appendix D, we report the results of OLS regressions that check the robustness of these results controlling for the FER score and demographic factors. These regression analyses confirm the findings reported here. Participants were significantly more likely to open the link in the email—and thus, start working on the real effort task to reveal their

FER test scores—in the *Immediate* treatment compared with those in the *Delayed* treatment, when we look at reasonably short time windows (any window within two hours), even after controlling for their performance and demographic factors.

This result suggests that demand for information is higher when an information gap is more top of mind. Although we cannot completely rule out that differential selection may play a role in this experiment, selection cannot fully account for this result. If participants in the *Delayed* treatment were busier at the time of receiving the email but still just as curious, then we would expect them to have greater demand for information than in the *Immediate* treatment in *later* time windows (when they eventually catch up on their email). However, we do not see any catch-up effect for participants in the *Delayed* treatment. The proportion of participants starting the extra task in later time windows is virtually identical across treatments (see Figure 6), even though they had *a week* after receiving the email to complete the task and obtain the information. This supports our view that higher salience leads to stronger demand for this information.

3.2.2.3. Experiment 2C: Valence. A total of 501 participants (94.5%) completed the experiment. We excluded three participants (0.6%) who submitted duplicate responses. The final sample contained 498 participants (55.2% female, $M_{age} = 37.3$ years): 246 in the *Easy* treatment and 252 in the *Hard* treatment.

3.2.2.3.1. Performance on the FER Test and Manipulation Check. Participants scored significantly higher in the *Easy* treatment, $M = 83.2\%$, than in the *Hard* treatment, $M = 47.4\%$, $t(478) = 39.814, p < 0.001$, Cohen's $d = 3.56$, 95% CI [33.99, 37.52]. Consistent with actual performances, people in the *Easy* treatment guessed that they scored significantly higher, $M = 69.3\%$, than participants in the *Hard* treatment, $M = 59.4\%$, $t(494) = 6.410$, $p < 0.001$, Cohen's $d = 0.57$, 95% CI [6.87, 12.94]. These differences in both actual and expected scores ensured that the majority of participants faced different scenarios in the two treatments. In the *Easy* treatment, 168 people (68.3%) had both scores above average and thus, could receive good news only; 73 (29.7%) could receive mixed news; and only 5 (2%) could receive bad news only. By contrast, in the *Hard* treatment, 163 (64.7%) could receive bad news only, 88 (34.9%) could receive either good or bad news, and only 1 person (0.4%) could receive good news only.

3.2.2.3.2. Main Results: Willingness to Exert Effort to Reveal Score. First, we compared the proportion of people who started and completed the extra task to reveal their score between experimental treatments. Consistent with our predictions, significantly more people started the extra task in the *Easy* treatment, $M = 63.4\%$, than in the *Hard* treatment, $M = 50.0\%$, $\chi^2(1, N = 498) = 8.583$, $p = 0.003$, Cohen's $w = 0.135$ (see Figure 7(a)). Similarly, significantly more people completed the task in the *Easy* treatment, $M = 58.5\%$, than in the *Hard* treatment, $M = 43.7\%$, $\chi^2(1, N = 498) = 10.450$, $p = 0.001$, Cohen's $w = 0.149$.

We also looked at the proportion of people who started and completed the extra task in both treatments, depending on whether they were facing good news only, bad news only, or mixed news. Although only 78 of 168 people (46.4%) started the extra task when facing bad news only, a marginally significantly larger proportion of people did so when facing mixed news, $M = 57.1\%$, $\chi^2(1, N = 329) = 3.362, p = 0.067$, Cohen's $w = 0.107$. The proportion of people starting the extra task was even higher among people who faced good news only: 112 of 169 participants (66.3%). This proportion is significantly higher than the proportion in the bad news only scenario, $\chi^2(1, N = 337) = 12.695$, $p < 0.001$, Cohen's $w = 0.200$, but not significantly higher

Figure 6. (Color online) Main Results, Experiment 2B

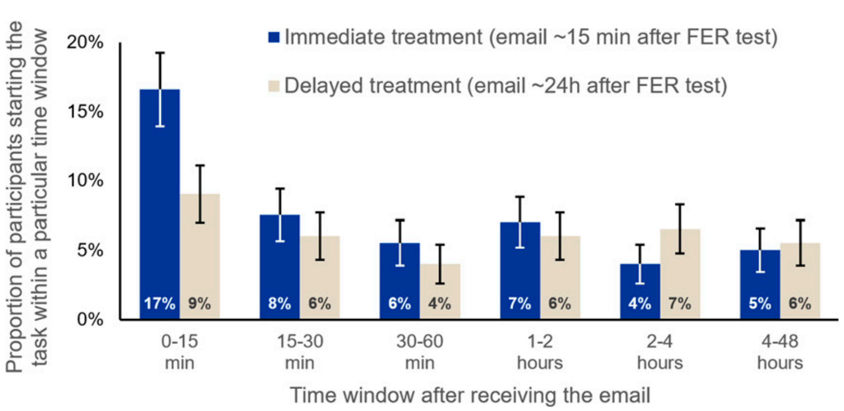

Note. Error bars represent ±1 standard error.

Figure 7. (Color online) Main Results and Mediation Analysis, Experiment 2C

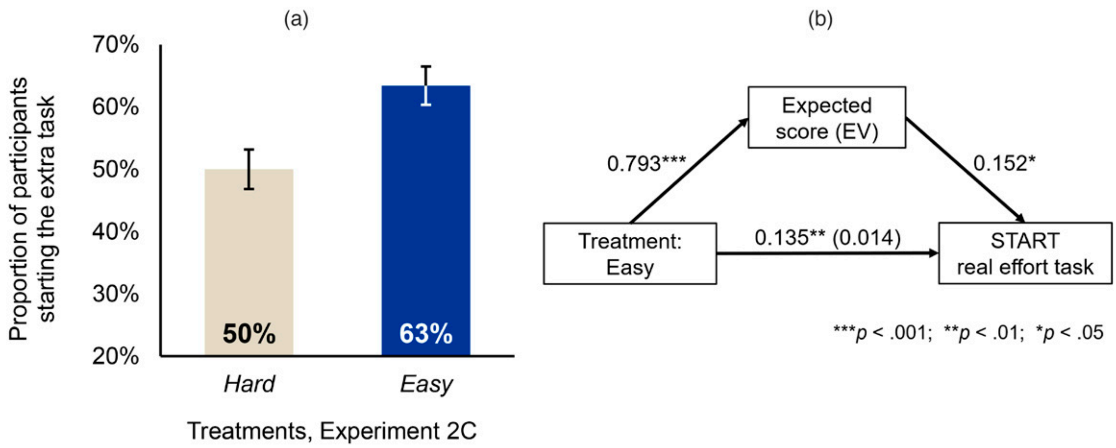

Notes. Error bars in panel (a) represent ±1 standard error. Coefficients in panel (b) are standardized β-coefficients. EV, expected value.

than in the mixed news scenario, $\chi^2(1, N = 330) = 2.537$, $p = 0.111$, Cohen's $w = 0.094$.

The results are similar if we look at the proportion of people completing the extra task, which was lowest among people who faced bad news, $M = 38.7\%$, followed by the mixed news scenario, $M = 54\%$, and it was highest in the good news scenario, $M = 60.4\%$. The proportion in the bad news scenario was significantly lower than in the good news scenario, $\chi^2(1, N = 337) = 14.964, p < 0.001$, Cohen's $w = 0.217$, and it was significantly lower than in the mixed news scenario, $\chi^2(1, N = 329) = 7.185$, $p = 0.007$, Cohen's $w = 0.154$. The proportion in the mixed news scenario was not significantly different from the proportion in the good news scenario, $\chi^2(1, N = 330) = 1.099$, $p = 0.295$, Cohen's $w = 0.064$.

These results are confirmed by OLS regression analyses that control for the total time spent on the FER test and demographics. In these analyses, we also tested whether the expected score (i.e., the average of the two possible scores) predicts participants' willingness to start and complete the extra task (see Table D5 in Online Appendix D). Our theory predicts that people are more motivated to obtain information if they are expecting to learn good news (i.e., a positive relationship between expected score and willingness to reveal the actual score). Consistent with our theory, we found that expected score significantly predicts both the willingness to start and the willingness to complete the extra task, $\beta = 0.003$, $SE = 0.002$, $t(495) = 2.094$, $p = 0.037$ and $\beta = 0.004$, $SE = 0.002$, $t(495) = 2.636$, $p = 0.009$ for starting and completing, respectively) (see Table D5 in Online Appendix D). Moreover, including expected score as a predictor makes the treatment dummy variable nonsignificant, both $p > 0.842$, which suggests that being assigned to the *Easy* or the *Hard* treatment only affects willingness to reveal the actual score through its effect on the expected score.

3.2.2.3.3. Mediation Analyses. To test whether expected score mediates the effect of experimental manipulation on the willingness to start and complete the effort task, we conducted a mediation analysis. We included the experimental condition as the predictor variable and starting and completing the extra task as outcome variables. We added the expected score as the proposed mediator variable. All variables were standardized before conducting the mediation analysis. A bootstrapped mediation with 5000 replications revealed that the expected score *fully* mediates the effect of the experimental condition on starting the extra task: $\beta = 0.121$, 95% CI [0.006, 0.233], $p = 0.041$ (see Figure 7(b)). The results are similar when looking at the effect of experimental condition on completing the extra task: $\beta = 0.152$, 95% CI [0.039, 0.264], $p = 0.010$.

3.2.3. Discussion. The results of these experiments provide additional support for Hypotheses 1 and 2, showing that, in an ego-relevant domain, higher perceived importance and higher salience increase demand for information (Experiments 2A and 2B, respectively). The ego-relevant domain also allows us to manipulate the valence of anticipated beliefs, confirming the prediction that demand for information is higher when the valence of expected news is more positive (Hypothesis 3).

4. General Discussion

4.1. Additional Predictions About Curiosity

Our experiments confirm our theory's predictions that demand for information can be stimulated or inhibited by manipulating the importance of a question, the salience of a question, or the valence of the potential answers to a question. We also derive additional predictions that we do not test here but that make sense of existing findings about curiosity.

Our theory predicts that another attentional factor, (recent) *surprise*, can stimulate curiosity too. If a person has previously received information addressing (but not completely resolving) a question and has not yet adapted to his revised beliefs, curiosity about this question tends to be stronger than if no relevant information had been received (or than if adaptation had already occurred) and even stronger when the previously received information was more surprising. Although we do not test for an effect of surprise in our experiments, this prediction is in line with existing evidence that being surprised stimulates curiosity about trivia questions (Loewenstein 1994, Vogl et al. 2020, Dubey et al. 2021). For example, people are more curious to find out the easternmost state in the United States (which surprisingly, happens to be Alaska) after getting separate feedback on each of three wrong guesses they made than after getting feedback all at once on three wrong guesses (Loewenstein 1994). Similarly, incrementally revealing hidden information about the identity of the protagonist in a story keeps people more curious throughout the story (Law et al. 2016), and incrementally revealing attributes of a vacation package makes people more curious about it (Wright et al. 2018). The provision of each piece of information generates surprise and increases curiosity.

According to our theory, curiosity also depends on the expected *informativeness* of a piece of information but only to the extent the information addresses specific questions attracting attention. People are more curious when they expect that information will more completely resolve an information gap, and people are especially motivated to acquire information that has the potential to fill multiple information gaps at once. Simultaneously resolving multiple information gaps generates an *epiphany*—a eureka moment of sudden comprehension. People may be especially curious when they anticipate a potential epiphany. Future work could test this prediction.

4.2. Implications for Information Avoidance

Our theory predicts that noninstrumental information tends to be desired when it addresses an attention-grabbing question (i.e., an information gap), as long as a person anticipates nonnegative beliefs (i.e., beliefs an individual does not mind thinking about). When filling an information gap poses no threat to utility, as would be true for answering a purely "intellectual" question (e.g., whether a particular tree is an oak or an elm), people generally want the information. However, when acquiring information might lead to negative beliefs, individuals may choose to avoid this information.

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. As the valence of anticipated outcomes becomes more negative, information avoidance becomes stronger (e.g., Ganguly and Tasoff 2016, Charpentier et al. 2018). The Bénabou and Tirole (2002) model of self-confidence and the Kőszegi (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 people with positive beliefs want to hold on to their current beliefs, whereas people with negative beliefs want to revise them.[16] Although the logic is intuitive, the empirical research suggests that this is not typically the case.

According to our theory, preference about information that would more clearly resolve negative beliefs involves a trade-off between curiosity and the desire to not think about these negative beliefs. We predict that this trade-off may depend on the prior attention directed toward these beliefs (as well as how negatively valenced they are). If the marginal increase in attention because of surprise is independent of the salience and importance, then 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. Indeed, van Dijk and Zeelenberg (2007) and Falk and Zimmermann (2016) manipulate salience (in different ways) to affect willingness to obtain potentially negative information, in line with our prediction. Along these lines, many people may avoid medical tests to avoid thinking about the possibility of being sick, but when forced to reckon with it (e.g., when talking with a doctor about symptoms that cannot be ignored), they may then prefer to be informed of a diagnosis. In Online Appendix E, we use our theory to provide an alternative account of avoidance of medical testing in the context of genetic testing for Huntington's Disease (Oster et al. 2013).

4.3. The Belief Resolution Effect

According to our theory, the impact of new information on attention is greatest when uncertainty about a question is resolved completely. Surprise prompts an immediate spike in attention, but it fades with adaptation. The underlying question then becomes unimportant because with the answer known, there is no longer a range of possible answers. The *belief resolution effect* refers to the dynamic pattern of attention that results from filling an information gap and then adapting to it. When an answer is learned with certainty, there is an immediate boost in attention weight on it, but after the person adapts, this attention weight falls.

A surprising feature of curiosity discussed in the review of Loewenstein (1994) 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 but soon forgotten. This property is naturally accommodated by the belief resolution effect. The attention weight associated with a particular question initially rises when the definitive answer is learned but ultimately falls below its prior level after a person adapts. 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.

The belief resolution effect also implies that the ostrich effect for unpleasant information may be counterproductive to individual welfare. Although people may avoid bad news because they do not want to think about it, the effect on attention is likely to reverse after people adapt. According to the belief resolution effect, after people adapt to new definitive beliefs, surprise fades, and certainty allows one to pay less attention to bad news, as it eventually seems less important. This can facilitate hedonic adaptation (e.g., Smith et al. 2009). So, it might be better initially to have definitive good news and worse to have definitive bad news, but eventually, the situation is likely to change because people adapt to both good and bad news, when it is definitive. Although 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. Our theory helps to explain why many people avoid confronting issues they do not like thinking about, and it also predicts that people with greater foresight will be more likely to choose to obtain information about such issues.

The same situation, but in reverse, occurs for positive information. Uncertainty can prolong the pleasure of good news. Wilson et al. (2005) 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, although people were unaware that this would be 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 adaptation to bad news, 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" and 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). Additionally, Ho et al. (2021) observe a significant correlation between discount rates and information avoidance across a variety of domains. 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.

5. Conclusion

In this paper, we use a theory of utility from beliefs about information gaps to make sense of a wide range of phenomena involving the demand for, or in some cases, the desire to avoid obtaining, information. The theory can be applied to understand the effectiveness of clickbait—headlines that raise salient questions and promise answers for those who click on them (Blom and Hansen 2015, Venneti and Alam 2018)—the backlash to mandatory disclosure of calorie information (Loewenstein et al. 2014), and the avoidance of medical tests (Thornton 2008, Hertwig and Engel 2020). The standard account of the economics of information, which assumes that information is desired only 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 ignore bad news. Economists have addressed some of these motives in isolation (e.g., Caplin and Leahy 2001, Bénabou and Tirole 2002, Brunnermeier and Parker 2005, Kőszegi 2006, Bénabou 2013), but the theory developed and tested here integrates a wide range of these motives into a unified theory. Sharot and Sunstein (2020) also have offered a framework for fitting these motives together; our theory fits into their framework but relies on more specific assumptions, and thus, it makes specific new testable predictions, which we show to be empirically supported.

Although our modeling relies on an extensive new apparatus, including the concepts of questions, answers, and attention weights, it offers many new predictions. As detailed in Section 2.4, we predict that contextual factors that affect attention, such as the importance and the salience of an information gap, will affect the demand for (and avoidance of) information, as will the valence of potential beliefs. We find support for these predictions in our experiments. Consistent with existing empirical evidence, but not tested here, the model also predicts that providing some related information (especially if surprising) increases demand for information and that there is greater demand for information that may fill more information gaps, as long as only nonnegative beliefs are expected. In addition to these derived predictions, we can also identify some additional predictions that go beyond our formal modeling but that follow conceptually from our underlying theory. First, individuals who anticipate adaptation and who discount the future less should be more likely to expose themselves to information relating to negative beliefs and less eager to obtain information relating to positive beliefs. Second, anticipation that receipt of information will occur, especially in a context that makes relevant questions 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 self-reports, facial coding, physiological measurements, or even brain activity scans (see, e.g., Ruan et al. 2018). First, 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. Second, 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. Third, if one can anticipate that a latent, meaningful question has nonnegative valence answers, then activating the question and learning the answer leave one better off than not being aware of the question in the first place (Golman and Loewenstein 2018b).

Our framework can help to shed light not only on information acquisition and avoidance but on other phenomena as well. In a companion paper (Golman et al. 2021), we argue that attention-based utility arising from information gaps also underlies an alternative account of risk and ambiguity preferences that is conceptually different from, and has different testable implications from, the usual account of risk preferences involving utility curvature 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.

The question of when people seek out or avoid information has gained importance in the internet age, with so much information available at our fingertips. Attention has become a highly valued and sought commodity (Simon 1971, Davenport and Beck 2001). Competing for consumers' attention, media organizations and digital marketing professionals have become ever more clever about creating clickbait that opens information gaps and piques curiosity. Television producers have likewise mastered the art of ending episodes with "cliffhangers" that open information gaps and beckon the viewer on to the next episode in search of answers. Apps aiming to help patients manage health conditions and investors manage their finances seek to engage curiosity to overcome information avoidance (because such information is, inevitably, sometimes adverse). As these content creators, product developers, public relations managers, and marketers all vie for consumers' attention and as consumers must sift through (and sometimes resist) their appeals and policy makers must figure out how to promote legitimate information, they all can benefit from a better understanding of the theoretical underpinnings of information seeking and avoidance.

Endnotes

[1] These motives fit into the Sharot and Sunstein (2020) framework as well. Motivated attention to favorable beliefs (and away from unfavorable beliefs) has "hedonic value," and curiosity has "cognitive value."

[2] Karlsson et al. (2009) escapes this critique by assuming that beliefs have less impact on utility in the absence of information, and Brunnermeier and Parker (2005) and Oster et al. (2013) get around it by relaxing Bayesian updating. The Bénabou and Tirole (2002) and Kőszegi (2006) models predict information avoidance specifically when people have *high* self-confidence.

[3] Oxford Languages defines *important* as "of great significance or value; likely to have a profound effect on success, survival, or well-being." Our characterization aligns with this definition, invoking the concept of utility in place of "success, survival, or well-being." Our operationalization reflects both the likelihood and potential magnitude of a possible effect on utility.

[4] Nonlinear probability weighting or reference-dependent valuation is as plausible for information preferences as for those involving only outcomes but would add extra complexity.

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

[6] We admit that people do not typically accurately assess the usefulness of information (see, e.g., Hoffman 2016). For simplicity, we assume that people know their own utilities, but this assumption could be modified to allow for heuristic assessment of instrumental value.

[7] Proposition 1 in Online Appendix A formally derives these predictions. The validity of Hypotheses 1 and 2 in the domain of negative beliefs relies on an ancillary assumption that the effect of surprise on attention is independent of the prior level of attention; these predictions about decreased information avoidance could fail if prior attention amplifies surprise, but they would still hold as predictions about demand for information in the domain of neutral or positive beliefs.

[8] The *Delayed* treatment could also decrease the importance of the information gap—as time passes, participants may adapt to the fact that they did not win the bonus, and they may then feel that the answer to the puzzle matters less to them. Still, the primary pathway for this manipulation to affect attention is likely that participants get distracted by other thoughts in the intervening 24 hours (i.e., through decreased salience).

[9] We also conducted two pilot experiments prior to Experiment 1A; see the online appendix. These pilots had mostly identical experimental designs to those of Experiment 1A, with some trivial differences (e.g., max number of clicks was 5, not 10, and participants were presented all puzzles on the same screen instead of one by one). Both pilot experiments tested Hypothesis 1 importance and yielded significant results consistent with Hypothesis 1; the effect sizes are comparable with those observed in Experiment 1A.

[10] Although not crucial for testing our theoretical predictions, we also analyzed the behavior of participants who correctly solved this puzzle. Unsurprisingly, they had little motivation to reveal the solution—that they already knew—and we do not find any significant effect of the experimental manipulation on their behavior. We report detailed test statistics in Online Appendix D.

[11] We report the significance value extracted from a Fisher's exact test because the low expected frequencies violate the assumptions necessary for the chi-squared test. The corresponding chi-square statistic with Yates correction would be 3.579 ($p = 0.059$).

[12] Similar but not identical tests are available on the internet, but they would have required participants to take another test all over again (which would take longer than three minutes). Thus, only we could reveal their score on the test they already took.

[13] Technically, SD(U) is a measure of the importance of winning the bonus, not a direct measure of the importance of their score, because the dependency is just on whether they win the bonus rather than on each possible score they could have gotten. However, eliciting this measure was much less burdensome on participants, and we assume that feeling that their score is more important is highly correlated with feeling that winning the bonus is more important.

[14] This was randomized across participants; within both the *Easy* and *Hard* treatments, 50% of participants had an alternative score that was 20 points higher than their actual score, and 50% of participants had an alternative score that was 20 points lower than their actual score. If an alternative score would have been lower than 0% or higher than 100% as the result of the calculation, we adjusted it to 0% or 100%, respectively. For those people who scored above 97% ($n = 8$), we always applied the –20 point adjustment to avoid having both scores in the same (21st) score bin.

[15] As a further robustness check, we control for the fixed pay amount in the regression analyses; see Table D3 in Online Appendix D.

[16] If people instead had concave belief-based utility for bad news and convex belief-based utility for good news, they would have greater demand for information when they had more positive beliefs, but this specification of reference-dependent belief-based utility relies on the opposite curvature assumptions of standard prospect theory.

References

Alos-Ferrer C, Garcia-Segarra J, Ritschel A (2018) Performance curiosity. *J. Econom. Psych.* 64:1–17.

Andries M, Haddad V (2020) Information aversion. *J. Political Econom.* 128(5):1901–1939.

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.

Bénabou R (2013) Groupthink: Collective delusions in organizations and markets. *Rev. Econom. Stud.* 80(2):429–462.

Bénabou R, Tirole J (2002) Self-confidence and personal motivation. *Quart. J. Econom.* 117(3):871–915.

Berlyne DE (1954) A theory of human curiosity. *British J. Psych.* 45 (3):180–191.

Blom JN, Hansen KR (2015) Click bait: Forward-reference as lure in online news headlines. *J. Pragmatics* 76:87–100.

Brunnermeier M, Parker J (2005) Optimal expectations. *Amer. Econom. Rev.* 95(4):1092–1118.

Burks S, Carpenter J, Goette L, Rustichini A (2013) Overconfidence and social signalling. *Rev. Econom. Stud.* 80(3):949–983.

Buyalskaya A, Camerer CF (2020) The neuroeconomics of epistemic curiosity. *Current Opinion Behav. Sci.* 35:141–149.

Cabrales A, Gossner O, Serrano R (2013) Entropy and the value of information for investors. *Amer. Econom. Rev.* 103(1):360–377.

Camerer C, Loewenstein G, Weber M (1989) The curse of knowledge in economic settings: An experimental analysis. *J. Political Econom.* 97(5):1232–1254.

Caplin A, Leahy J (2001) Psychological expected utility theory and anticipatory feelings. *Quart. J. Econom.* 116(1):55–79.

Charpentier CJ, Bromberg-Martin ES, Sharot T (2018) Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proc. Natl. Acad. Sci. USA* 115(31):E7255–E7264.

Chater N, Loewenstein G (2015) The underappreciated drive for sense-making. *J. Econom. Behav. Organ.* 126:137–154.

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

Dana J, Weber R, Kuang J (2007) Exploiting moral wiggle room: Experiments demonstrating an illusory preference for fairness. *Econom. Theory* 33(1):67–80.

Davenport T, Beck J (2001) *The Attention Economy: Understanding the New Currency of Business* (Harvard Business School Press, Cambridge, MA).

Dillenberger D (2010) Preferences for one-shot resolution of uncertainty and allais-type behavior. *Econometrica* 78(6):1973–2004.

Dubey R, Mehta H, Lombrozo T (2021) Curiosity is contagious: A social influence intervention to induce curiosity. *Cognitive Sci.* 45(2):e12937.

Eil D, Rao J (2011) The good news-bad news effect: Asymmetric processing of objective information about yourself. *Amer. Econom. J. Microeconomics* 3(2):114–138.

Eliaz K, Schotter A (2010) Paying for confidence: An experimental study of the demand for non-instrumental information. *Games Econom. Behav.* 70(2):304–324.

Eliaz K, Spiegler R (2006) Can anticipatory feelings explain anomalous choices of information sources? *Games Econom. Behav.* 56(1):87–104.

Ely J, Frankel A, Kamenica E (2015) Suspense and surprise. *J. Political Econom.* 123(1):215–260.

Falk A, Zimmermann F (2016) Beliefs and utility: Experimental evidence on preferences for information. Working Paper No. 6061, CESifo, Munich.

Faul F, Erdfelder E, Buchner A, Lang A-G (2009) Statistical power analyses using G*Power 3.1: Tests for correlation and regression analyses. *Behav. Res. Methods* 41(4):1149–1160.

Gabaix X, Laibson D, Moloche G, Weinberg S (2006) Costly information acquisition: Experimental analysis of a boundedly rational model. *Amer. Econom. Rev.* 96(4):1043–1068.

Ganguly A, Tasoff J (2016) Fantasy and dread: The demand for information and the consumption utility of the future. *Management Sci.* 63(12):4037–4060.

Gigerenzer G, Garcia-Retamero R (2017) Cassandra's regret: The psychology of not wanting to know. *Psych. Rev.* 124(2):179–196.

Golman R, Loewenstein G (2018a) Information gaps: A theory of preferences regarding the presence and absence of information. *Decision* 5(3):143–164.

Golman R, Loewenstein G (2018b) The desire for knowledge and wisdom. Gordon G, ed. *The New Science of Curiosity* (Nova, Hauppauge, NY), 37–42.

Golman R, Gurney N, Loewenstein G (2021) Information gaps for risk and ambiguity. *Psych. Rev.* 128(1):86–103.

Golman R, Hagmann D, Loewenstein G (2017) Information avoidance. *J. Econom. Literature* 55(1):96–135.

Gottlieb J, Hayhoe M, Hikosaka O, Rangel A (2014) Attention, reward, and information seeking. *J. Neuroscience* 34(46):15497–15504.

Gottlieb J, Oudeyer PY, Lopes M, Baranes A (2013) Information seeking, curiosity and attention: Computational and empirical mechanisms. *Trends Cognitive Sci.* 17(11):585–593.

Grant S, Kajii A, Polak B (1998) Intrinsic preference for information. *J. Econom. Theory* 83(2):233–259.

Hertwig R, Engel C (2016) Homo ignorans: Deliberately choosing not to know. *Perspect. Psych. Sci.* 11(3):359–372.

Hertwig R, Engel C (2020) *Deliberate Ignorance* (MIT Press, Cambridge, MA).

Hirsh J, Inzlicht M (2008) The devil you know: Neuroticism predicts neural response to uncertainty. *Psych. Sci.* 19(10):962–967.

Hirshleifer J, Riley J (1979) The analytics of uncertainty and information—an expository survey. *J. Econom. Literature* 17 (4):1375–1421.

Ho E, Hagmann D, Loewenstein G (2021) Measuring information preferences. *Management Sci.* 67(1):126–145.

Hoffman M (2016) How is information (under-) valued? Evidence from framed field experiments. *Econom. J.* 126(595):1884–1911.

Hsee CK, Ruan B (2016) The pandora effect: The power and peril of curiosity. *Psych. Sci.* 27(5):659–666.

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

Kang MJ, 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. *Psych. Sci.* 20(8):963–973.

Karlsson N, Loewenstein G, Seppi D (2009) The ostrich effect: Selective attention to information. *J. Risk Uncertainity* 38(2):95–115.

Kidd C, Hayden B (2015) The psychology and neuroscience of curiosity. *Neuron* 88(3):449–460.

Kőszegi B (2006) Ego utility, overconfidence, and task choice. *J. Eur. Econom. Assoc.* 4(4):673–707.

Kőszegi B (2010) Utility from anticipation and personal equilibrium. *Econom. Theory* 44(3):415–444.

Kreps D, Porteus E (1978) Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46(1):185–200.

Kruger J, Evans M (2009) The paradox of alypius and the pursuit of unwanted information. *J. Experiment. Soc. Psych.* 45(6):1173–1179.

Kupor D, Tormala Z (2015) Persuasion, interrupted: The effect of momentary interruptions on message processing and persuasion. *J. Consumer Res.* 42(2):300–315.

Langner O, Dotsch R, Bijlstra G, Wigboldus DHJ, Hawk ST, van Knippenberg A (2010) Presentation and validation of the Radboud Faces Database. *Cognition Emotion* 24(8):1377–1388.

Law E, Yin M, Goh J, Chen K, Terry MA, Gajos KZ (2016) Curiosity killed the cat, but makes crowdwork better. *Proc. 2016 CHI Conf. Human Factors Comput. Systems* (ACM, San Jose, CA), 4098–4110.

Liquin EG, Lombrozo T (2020) A functional approach to explanation-seeking curiosity. *Cognitive Psych.* 119:101276.

Loewenstein G (1994) The psychology of curiosity: A review and reinterpretation. *Psych. Bull.* 116(1):75–98.

Loewenstein G, Sunstein CR, Golman R (2014) Disclosure: Psychology changes everything. *Annual Rev. Econom.* 6:391–419.

Lucey P, Cohn JF, Kanade T, Saragih J, Ambadar Z, Matthews I (2010) The Extended Cohn-Kanade Data Set (CK+): A complete data set for action unit and emotion-specified expression. *2010 IEEE Comput. Soc. Conf. Comput. Vision Pattern Recognition Workshops*, 94–101.

Ma DS, Correll J, Wittenbrink B (2015) The Chicago face database: A free stimulus set of faces and norming data. *Behav. Res. Methods* 47(4):1122–1135.

Möbius M, Niederle M, Niehaus P, Rosenblat T (2022) Managing self-confidence: Theory and experimental evidence. *Management Sci.* Forthcoming.

Olafsson A, Pagel M (2017) The ostrich in us: Selective attention to financial accounts, income, spending, and liquidity. NBER Working Paper No. 23945, National Bureau of Economic Research, Cambridge, MA.

Oster E, Shoulson I, Dorsey ER (2013) Optimal expectations and limited medical testing: Evidence from Huntington disease. *Amer. Econom. Rev.* 103(2):804–830.

Pagel M (2018) A news-utility theory for inattention and delegation in portfolio choice. *Econometrica* 86(2):491–522.

Picone G, Sloan F, Taylor D (2004) Effects of risk and time preference and expected longevity on demand for medical tests. *J. Risk Uncertainity* 28(1):39–53.

Ruan B, Hsee CK, Lu ZY (2018) The teasing effect: An underappreciated benefit of creating and resolving an uncertainty. *J. Marketing Res.* 55(4):556–570.

Shannon C (1948) A mathematical theory of communication. *Bell System Tech. J.* 27(3):379–423.

Sharot T, Sunstein CR (2020) How people decide what they want to know. *Nature Human Behav.* 4(1):14–19.

Sicherman N, Loewenstein G, Seppi D, Utkus S (2015) Financial attention. *Rev. Financial Stud.* 29(4):863–897.

Simon HA (1971) *Designing Organizations for an Information-Rich World* (Johns Hopkins University Press, Baltimore, MD).

Smith DM, Loewenstein G, Jankovich A, Ubel PA (2009) Happily hopeless: Adaptation to a permanent, but not to a temporary, disability. *Health Psych.* 28(6):787–791.

Stigler G (1961) The economics of information. *J. Political Econom.* 69(3):213–225.

Thornton R (2008) The demand for, and impact of, learning HIV status. *Amer. Econom. Rev.* 98(5):1829–1863.

van Dijk E, Zeelenberg M (2007) When curiosity killed regret: Avoiding or seeking the unknown in decision-making under uncertainty. *J. Experiment. Soc. Psych.* 43(4):656–662.

Venneti L, Alam A (2018) How curiosity can be modeled for a clickbait detector. Preprint, submitted June 11, https://arxiv.org/abs/1806.04212v1.

Vogl E, Pekrun R, Murayama K, Loderer K (2020) Surprised-curious-confused: Epistemic emotions and knowledge exploration. *Emotion* 20(4):625–641.

Wakker P (1988) Nonexpected utility as aversion of information. *J. Behav. Decision Making* 1(3):169–175.

Wilson T, Centerbar D, Kermer D, Gilbert D (2005) The pleasures of uncertainty: Prolonging positive moods in ways people do not anticipate. *J. Personality Soc. Psych.* 88(1):5–21.

Woolley K, Risen JL (2018) Closing your eyes to follow your heart: Avoiding information to protect a strong intuitive preference. *J. Personality Soc. Psych.* 114(2):230–245.

Wright S, Clarkson J, Kardes F (2018) Circumventing resistance to novel information: Piquing curiosity through strategic information revelation. *J. Experiment. Soc. Psych.* 76:81–87.