

## **PERSPECTIVE**

### **HOW PEOPLE DECIDE WHAT THEY WANT TO KNOW**

Tali Sharot<sup>1\*</sup> & Cass R. Sunstein<sup>2</sup>

<sup>1</sup>Affective Brain Lab, Department of Experimental Psychology, University College London, London, UK

<sup>2</sup>Harvard Law School, Harvard University, Cambridge, MA, USA

\* To whom corresponds should be addressed [t.sharot@ucl.ac.uk](mailto:t.sharot@ucl.ac.uk)

## **Abstract**

Immense amounts of information are now accessible to people, including information that bears on their past, present and future. An important research challenge is to determine how people decide to seek or avoid information. Here, we propose a framework of information-seeking that aims to integrate the motives that drive information-seeking and its avoidance. Our framework rests on the idea that information can alter people's *action*, *affect* and *cognition* in both positive and negative ways. The suggestion is that people assess these influences and integrate them into a calculation of the value of information that leads to information-seeking or avoidance. The theory offers a framework for characterizing and quantifying individual differences in information-seeking, which we hypothesize may also be diagnostic of mental health. We consider biases that can lead to both insufficient and excessive information-seeking and discuss how the framework offered can help government agencies to assess the welfare effects of mandatory information disclosure.

How do people decide what they want to know? As massive amounts of information are becoming available to people, this question is more pertinent today than ever. People must now decide whether they want to receive highly personalized information about their health, genetic make-up, and financial prospects, as well as those of their children. At the same time, new laws and regulations are being issued that require disclosure of information in diverse areas, including calories in food at chain restaurants, fuel economy of motor vehicles, energy efficiency of refrigerators, and genetic modification of food.

An important research challenge is to determine how people decide to seek or avoid information and how those decisions affect their welfare. By “information-seeking” we mean the active pursuit of knowledge, for example by asking questions, reading, running tests and conducting online searches. Understanding how people make such decisions is important for many fields, including the education and health sectors, for public policy and law, and for the design of intrinsically motivated, curious, artificial intelligence systems. Research on information-seeking has been surprisingly limited in comparison to other domains of human cognition and behaviour and somewhat narrow in focus, but seems to be experiencing revitalization in recent years (for review see 1). Below we present a framework for considering the motives that lead people to seek or avoid information, and explore how this framework may be used to investigate individual differences and facilitate policy making.

### **An Integrative Framework of Information-Seeking Motives**

Deciding whether to seek information is a particularly difficult problem to solve, because knowledge and its avoidance can serve diverse and sometimes competing functions (2). For instance, imagine we had information on whether you had a genetic predisposition to specific forms of cancer; would you want to know? When we posed this question to 400 individuals online just over half (58%) said yes and the rest (42%) said no. What drove people’s decisions?

Classic theories of information-seeking suggest that agents seek information that can aid decisions to obtain reward and avoid harm (i.e., that has *‘instrumental utility’*) (3). For example, knowing whether one carries the BRCA gene, which increases the likelihood of breast cancer, can inform decisions on whether to undergo preventative surgery.

It is apparent, however, that instrumental utility is not the sole factor guiding information-seeking, as agents often want information that cannot be used to alter outcomes (4-7). People might want information about the life of William Shakespeare or the origins of the universe even if that information will have no effect on what they do. This observation has led to the idea that people use a heuristic according to which “knowledge is always valuable” (4). Such an approach may be adaptive, because information could turn out to be useful in the future even if it appears useless at present (8-9). Consistent with this claim, neuroscientists have shown that the opportunity to gain knowledge is encoded using some of the same neural architecture and algorithms as for primary rewards (5-6, 10-19), suggesting that knowledge may have intrinsic value (4,7) (**Box 1**). ‘Information Prediction Error’ signals (IPEs) have been identified in dopamine rich brain regions (6) (**Box 1**), which analogous to Reward Prediction Errors (RPEs) (20), are theorized to provide reinforcement for seeking information (6).

Human behaviour, however, is inconsistent with the proposition that knowledge is always perceived to be valuable, as people sometimes choose to remain ignorant. For example, they

often reject medical screening (21-23). Examples of this kind have led psychologists, economists and neuroscientists to consider the hedonics of information (2,5,24-28). It is theorized that when deciding whether to seek information, people consider how knowledge will influence their psychological well-being (i.e. ‘how will I feel if I learn I carry the BRCA gene?’) and that, all else being equal, agents will be more likely to seek information when they expect good news than when they expect bad news (2,5,24-26).

It remains unclear, however, how instrumental utility and hedonic utility are integrated in the brain to guide information-seeking choices (see **Box 2** for outstanding questions and for a computational proposal see 29). It also remains unclear which other factors play a role. Here we propose a framework of information-seeking that aims to encompass the *motives* that drive information-seeking and its avoidance. This theoretical framework also offers a structure for characterizing and quantifying individual differences in information-seeking.

Our theory rests on the idea that information can alter people’s action, affect and cognition in both positive and negative ways (**Fig 1a**). When deciding whether to seek information people may estimate the expected impact of information on their **action** (will the knowledge help, hinder or have no influence on my ability to make decisions to increase reward and avoid harm?), **affect** (will the information induce positive or negative feelings, or have no influence on my affect?) and **cognition** (will information improve the ability to comprehend and anticipate my reality?). Each of these estimates can be positive (increasing information seeking), negative (increasing information avoidance) or zero (inducing indifference). These estimates will then be integrated into a computation of the value of information, which will trigger information seeking (if the integrated value is sufficiently positive), its active avoidance (if the integrated value is sufficiently negative), or neither (i.e. indifference). Each factor can be weighted differently, influencing the decision to seek/avoid information to different degrees (**Fig 1a**). Below we expand on each of these three factors.

- 1. Action (instrumental value).** The ability to *use* information to select actions that increase extrinsic rewards and help evade losses is an important driver of information-seeking. This component of our framework is found in most classic models of information-seeking (3). What has often been overlooked, however, is that information can also have *negative* instrumental value. That is, knowledge can at times cause individuals to select actions that lead to worse outcomes, while deliberate ignorance can lead to better outcomes (30). For example, not knowing the gender of an applicant could improve hiring decisions by reducing bias; not knowing whether a client is guilty could improve a solicitor’s performance; not being aware of one’s own limitations increases confidence and motivation, which at times may improve functioning (31).

According to rational choice theory, information cannot have negative instrumental value, because it could always be disregarded after it is obtained. In practice, however, once a piece of information is known, its impact on action is often beyond an agent’s control and/or awareness and thus cannot be reversed. For example, once an applicant’s gender is known, one cannot undo this knowledge and reverse its unconscious impact on hiring decisions. In law, the point is familiar in the context of jury deliberations. If a judge directs a jury “not to consider” evidence that has been

wrongfully put before it (such as evidence of prior crimes), there is a real question whether the jury will or can follow that direction.

When knowledge is predicted to influence one's own action negatively, information will have negative instrumental value and vice versa when the influence is predicted to be positive. If knowledge is predicted not to influence action, instrumental value will be zero.

- 2. Affect (*hedonic value*).** Knowledge can induce both positive and negative affect. Knowing that one has a predisposition for certain cancers, for example, can generate sadness or fear (22). Thus, all else being equal, individuals are motivated to avoid information that induces negative affect and seek information that evokes positive affect – using information to regulate emotion (26). Consistent with this notion are observations that investors monitor their portfolio more frequently when they expect their worth has gone up rather than down (32); that some people refuse to receive results of medical tests they have taken (33); and that monkeys select to know in advance the size of reward they are about to receive (6, 18-19). Furthermore, it has been shown that humans are more likely to seek and pay for information when they expect good news (e.g., information indicating financial profit) and are willing to pay to avoid information when they expect bad news (e.g., information indicating financial loss) (5). Key subcortical regions in the reward circuitry – the Ventral Tegmental Area (VTA) and Substantia Nigra (SN) – selectively code for the opportunity to receive information that is expected to convey good news as if such information has utility in and of itself (5). The code takes the form of an IPE. Similar coding is not observed for the opportunity to gain information about losses, suggesting that the intrinsic utility of knowledge is modulated by valence (5).

We suspect, however, that hedonic utility considerations alone may at times cause individuals to seek information that is expected to induce negative affect. For example, uncertainty about bad news may trigger prolonged anxiety *that is worse* than the expected reaction to such information. Thus, similar to cases in which people select to undergo painful procedures now (such as dental work) rather than later in order to minimize dread (34), people may choose to uncover bad news now rather than remain in the dark.

The opposite may also be true. People may prefer not to know in advance the outcome of a sports game, the gender of their unborn baby, or the location of a vacation being planned by their spouse, because they want to preserve the positive feeling of surprise. Lack of knowledge can produce anxiety, but it can also produce delight.

The above examples all converge on the same principle: all else being equal people will choose to seek-information when the affective response to knowing at present time is anticipated to be more positive than to remaining ignorant. When the reverse is true people will choose to avoid information.

- 3. Cognition** (cognitive value). Information can enhance or reduce people's sense that they comprehend the world around them. In particular, information alters people's internal *mental models*. Mental models are representation of concepts (e.g., 'dog', 'Shakespeare', 'mom', 'alien', 'democracy', 'cancer', 'money', 'self') and the relationship among them, which are used to comprehend and anticipate reality (35,36). Certain concepts (sometimes referred to as 'nodes' of the model) may be part of one person's mental model but absent in another's. Moreover, certain concepts (e.g., 'dog') can be frequently activated within one person's model and inter-connected to other concepts ('love', 'daughter', 'fun', 'food', 'home') but not in another person's model.

We suggest that people will be more likely to want information relating to concepts that are frequently activated and highly inter-connected to other concepts in their mental models (e.g., 'self', 'human'). Again, individual differences exist. For example, a dog owner would be more interested to learn whether dogs are related to wolves than a non-dog owner. This information will have greater positive cognitive value for the former than the latter, despite the information having no clear practical use for either. This is because 'dog' is a central concept in the former's mental model. If information is very weakly related to people's mental model (e.g., knowing the middle name of someone passing you on the street) it will have zero 'cognitive value'.

It has been suggested that people strive to minimize the difference between their mental models and external reality (37). This can be achieved in two ways. The first is to improve existing models by seeking out information that will generate new connections among concepts, or strengthen connections that are suspected but of which one is *uncertain* (38,39). This approach will improve the fit of the model to reality by refining the model which will increase the ability to comprehend and anticipate reality. The second approach, less intuitive, is to avoid information that one suspects will weaken strong existing interconnected ties within the model (e.g., disconformity information,40-42). In this approach, people maintain a fit between internal representation and external reality not by actively changing the model but by actively changing the reality of which they are aware. Such avoidance can improve the *sense* of comprehension at present, because disconformity information can cause confusion and a need to rebuild large parts of the model. Thus, disconformity information may be assigned negative 'cognitive value', despite the fact that avoiding such information could create less accurate mental models.

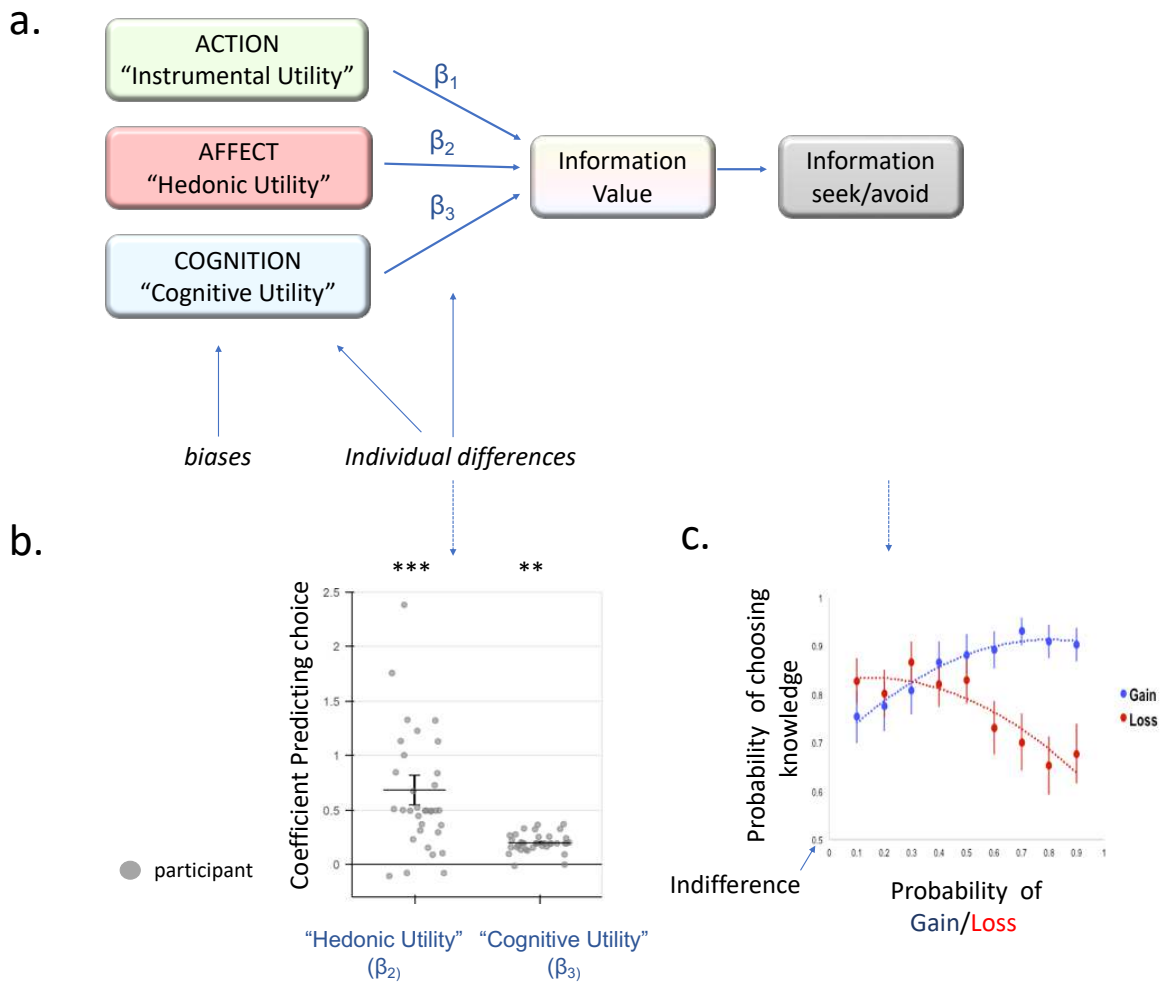
With the above we intend to outline a broad theoretical framework that can help guide future work in the behavioural sciences and neuroscience. The framework builds on many past studies, each highlighting different aspects of information-seeking motives. These include studies emphasizing the utility of beliefs (26,27), the need to "fill-in gaps of knowledge" (29) and the need for confirmation and confidence (41,42). It is also related to economic models outlining ways by which motives may alter information-seeking. For example, Golman et al., (28) highlight the importance of instrumental utility and suggest that information-seeking is further guided by motivated attention and curiosity. Note, that we view curiosity as an expression of information-seeking motive(s) rather than the primary motive itself. It is analogues to hunger, which is the feeling of needing food but not the reason food is needed

(the latter may be, for example, low sugar). We speculate that the feeling of curiosity does not simply map on to one of our three motives - a speculation that requires future testing (see also **Glossary** for a distinction between curiosity and information-seeking). With regards to motivated attention, we suspect that it is one of several ways affect may modulate information-seeking.

Our suggested framework could be used to explain information-seeking choices in different contexts and domains including health and finance, and to develop detailed models that predict such choices. We suggest that the three utilities could be measured and/or manipulated and that these measures could be used to explain and/or predict choice (note, however, that future work is required to identify the specific rules regarding which type of content under which circumstances generate which type of utility). For example, in a recent study it was found that participants' decisions to know or remain ignorant about the outcome of a lottery game was explained by hedonic and cognitive utility (in this task instrumental utility was set to zero by design) (5, **Figure 1b**). In particular, the more likely participants were to receive good news (about monetary gain) and the less likely they were to receive bad news (about monetary loss), the more they wanted to know their outcomes. Second, the more uncertain the outcome was, the more people wanted to know the outcome (5, **Figure 1c**). This information facilitated the ability to anticipate future events, but not to alter them. Similar results were observed in a separate study in which participants *could* use information to alter outcomes. In that study, instrumental utility was an additional orthogonal driver of information-seeking (43).

The aforementioned studies (5,43) were designed such that instrumental, hedonic and cognitive utilities fluctuated independently from each other. This made it possible to measure their separate effects on information-seeking. These motives can often be independent "in the wild." For example, a person may need to decide whether to undergo screening for a disease she is quite certain or uncertain she carries (impacting cognitive utility) that is either curable or not (impacting instrumental utility) such that the two dimensions are orthogonal. But in some cases the dimensions co-vary and can be difficult to disentangle. For example, all else being equal, the less curable a disease (instrumental utility) the worse it would be to learn that one carries it (hedonic utility).

Moreover, in the aforementioned studies (5,43) the expected influence of information on action, affect and cognition was quantified using objective cues that were available to participants. For example, the participants knew the exact probability of a gain/loss (5,43) or how useful a piece of information was for attaining rewards (43). Often, however, these quantities need to be estimated by the decision-maker. This aspect of the decision process, as we explain below, can result in information-seeking biases (**Figure 1a**).



**Figure 1. Integrative Framework of Information-Seeking Motives.** (a) Information seeking/avoidance is hypothesized to be driven by instrumental utility, hedonic utility and cognitive utility. These values reflect the predicted impact of information on action (will the knowledge help, hinder or have no influence on my ability to make decisions to increase reward and avoid harm?), affect (will the information induce positive or negative feelings, or have no influence on my affect?) and cognition (will information improve the ability to comprehend and anticipate my reality?), respectively. These estimates are integrated into a computation of the value of information, with different weights ( $\beta_{1-3}$ ) assigned to each of the three factors. The integrated value can be positive (increasing information seeking), negative (increasing information avoidance) or zero (leading to indifference). Biases in predicting the content of information and its influence on action, affect and cognition may result in suboptimal information-seeking behaviour. Individual differences in information-seeking are hypothesized to be related to variations in these estimations as well as differences in the weights ( $\beta_{1-3}$ ) assigned to the three drivers. (b) Figure adapted from (5) showing that information-seeking choices are explained by hedonic utility (in this case how likely is the information to be good/bad) and cognitive utility (in this case the amount of uncertainty which is reduced by information). In this study instrumental utility was set to zero. Each dot represents a participant. For illustration purposes, dots are scattered along the x-axis such that each dot is visible. (c) Findings from the same study (5) highlighting the impact of hedonic utility in a non-instrumental information-seeking task. Specifically, the more likely participants were to win a lottery the more they wanted to know the outcome, the more likely they were to loss, the less they wanted to know the outcome.

### Prediction Biases Influencing Information-Seeking



According to the above framework, the first stage in deciding whether to seek information involves solving a *prediction problem*. People must predict the likely content of information and its influence on action, affect and cognition. Thus, biases known to influence predictions will influence estimates of instrumental, hedonic and cognitive utility, altering information-seeking.

Prediction problems can be extremely hard to solve and notoriously vulnerable to biases. First, people have problems in estimating probability and thus may commit systematic errors in predicting the likely content of information (44,45). Second, there are potential gaps between the utility expected at the time of deciding whether to seek information and the utility experienced during consumption of information. Below are a few examples of such biases and their potential influence on information-seeking.

1. **Illusion of Control** (biasing predictions of instrumental utility). People tend to overestimate the ability of their actions to influence outcomes (46). This bias will likely lead to an over-estimation of the instrumental value of information. This is because individuals may believe that information about outcomes will enable them to take action to alter those outcomes, when in fact outcomes are beyond their control. This can lead people to overpay for useless information.
2. **Unrealistic Optimism** (biasing predictions of hedonic utility). People tend to overestimate the probability of positive outcomes and underestimate the probability of negative outcomes (44,47). When deciding whether to seek information, people may overestimate the probability of good news and underestimate the probability of bad news. This will lead to an overestimation of positive hedonic utility and an underestimation of negative hedonic utility – increasing information seeking. The result could be over-exposure to unexpected negative information and disappointment.
3. **Impact Bias** (biasing predictions of hedonic utility). People tend to overestimate the duration and intensity of future affective states (48). This bias could lead to an overestimation of positive hedonic utility (increasing information seeking for good news) and negative hedonic utility (increasing information avoidance for bad news). The latter can cause individuals to avoid negative information that could have otherwise had instrumental utility (including information about their health, finance and feedback about their performance), without the expected benefit to their emotional well-being.
4. **The illusion of knowledge** (biasing predictions of cognitive utility). People tend to underestimate how much new information will influence their understanding (49). Thus, people may underestimate the cognitive value of information.
5. **Overconfidence** (biasing predictions of cognitive utility). People tend to be overconfident in their knowledge (50). Thus people may underestimate how much new information may reduce uncertainty, biasing their estimate of its cognitive value.

The above is not meant to be an exhaustive list, but rather a demonstration of how prediction biases may lead to suboptimal information-seeking by influencing estimations of each of the three factors. People's motivation to seek information and their willingness to pay for information will capture the *expected* welfare effects from information, including instrumental, hedonic, and cognitive effects, but will often fail to accurately

capture the *experienced* welfare effects, thus leading to suboptimal information-seeking choices. The disjunction between the expected utility of information and the actual experienced utility of information will vary across individuals as a function of individuals' susceptibility to different biases.

### **Individual Differences in Information-Seeking: A Window into Mental Health?**

Individual differences in information-seeking (**Figure 1a,b**) may reflect differences in estimated hedonic, instrumental and cognitive utilities (which could be partially due to susceptibility to the biases described above) as well as differences in the weighting of those utilities (e.g. 5). We hypothesize that such differences may be related to psychiatric symptoms and conditions. The reason for this hypothesis is that many conditions, including depression and anxiety, are characterized by symptoms that could manifest as abnormalities in information-seeking (51-56). Anxiety, for example, is characterized by intolerance to uncertainty (57), which may lead to increased weighting of cognitive utility. Depression is characterized by reduced sensitivity to valence (58-60) and a reduction in the belief that one has control over outcomes (61), which can lead to a reduction in the impact of hedonic and instrumental utility on information-seeking, respectively.

It is likely that differences in information-seeking are not merely correlated with psychopathology symptoms, but rather that a bidirectional causal relationship exists. That is, psychiatric disorders may influence information-seeking behaviour and at the same time certain patterns of information-seeking may either protect individuals from psychiatric symptoms or make them more vulnerable. For example, if individuals fail to take into account the expected valence of knowledge, they may be overexposed to information that induces negative affect, thus increasing the likelihood of depressive symptoms.

It is interesting to consider whether certain psychotropic drugs reduce symptoms partially by altering such information-seeking patterns in patients. Dopamine function, in particular, is thought to play a role in information seeking, perhaps in modulating the impact of affect on information-seeking (**Box 1**). If that is true, conditions characterized by abnormal dopaminergic function (such as Parkinson's disease, schizophrenia and depression, 62-64) may be characterized by atypical information-seeking. Indeed, there have been some reports of atypical information-seeking in such patients (52,53). Moreover, prescription drugs targeting dopamine function may alter patients' information-seeking behavior with potential consequences, positive or negative, to patients' well-being.

An established relationship between information-seeking patterns and mental health would have important practical applications. Because people constantly engage in information-seeking online, a stream of data could theoretically be used to facilitate clinical diagnosis, monitoring, and treatment selection. Indeed, there is growing interest in using online behavior for diagnosis and stratification of psychiatric and neurological conditions, as data can be obtained frequently and non-intrusively. For example, recent work has begun to link keystroke patterns to motor disorders (65) and affective states (66). Such endeavours could potentially be boosted for diagnosis if rich behaviorally relevant features could be extracted, such as measured markers of information-seeking. Efforts in this vein should protect personal privacy.

### Policy Implications.

U.S. agencies are required to quantify the costs and benefits of information disclosure mandates (67). The requirement applies to a wide range of regulations, involving (for example) disclosure of fuel economy, graphic warnings for cigarettes, calorie labels, genetically modified foods, minerals used to finance mass atrocities, and country-of-origin labels (68). Our framework could be useful in helping officials do better in identifying the costs and benefits of disclosure. It could also trigger ideas of how to make information more attractive, thus increasing the likelihood that people will actually read leaflets and labels, and so benefit from them.

Currently, it is common for agencies to try to project the instrumental utility of disclosing (or not disclosing) information (69). For example, an agency might attempt to estimate the health benefits of disclosing the caloric contents of food in restaurants (70). The impact of information on non-instrumental utility, however, is often overlooked. This presents two problems. First, considering only instrumental utility does not give a complete picture of the actual costs and benefits. As we have seen, people might experience significant hedonic losses as well as gains as a result of receiving information or remaining ignorant (71).

Second, without considering the effect of non-instrumental utilities, policymakers will systematically mispredict the impact of information on people's action, because they will misjudge the likelihood that people will "consume" (i.e. read/ hear about) the information in the first place. For example, in estimating the health benefits of disclosing caloric content, policy-makers tend to assume that consumers generally will read those labels (70). Yet, if the information is likely to induce negative affect, the consumer may decide not to do so. Thus, the benefits of such labels will be overestimated. In fact, it has been shown that people select not to learn the caloric content of a tempting dessert, as they are aware this knowledge will reduce the likelihood they would select said dessert (72).

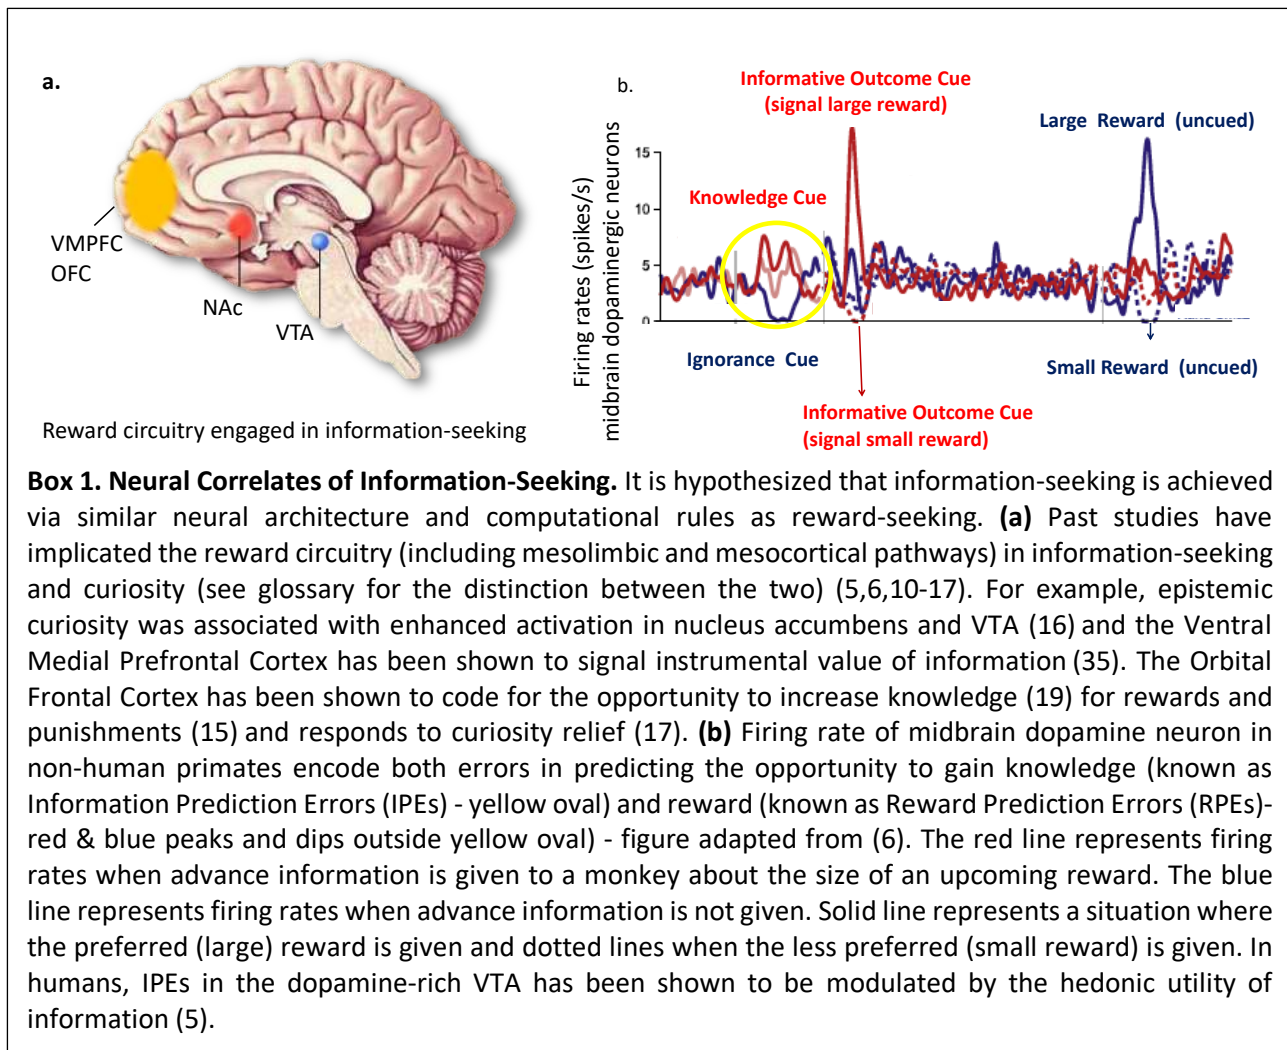
Considering the expected hedonic, cognitive and instrumental utility of information can also reveal how information should be framed in order to maximize use. That is, information should be framed in a way that highlights its ability to guide **action** and/or improve **affect** and/or **cognition**. For example, in order to encourage people to attend to pre-flight safety instructions, airlines have in recent years created videos that induce positive affect using humor, music and alluring images of vacation destination. By doing so, they overcame people's tendency to avoid such information in order to suppress flying anxieties. This approach has been so successful that millions now watch these videos even at home (73).

Finally, regulators should consider individual differences in information-seeking patterns and the influence of information on welfare due to mental health issues (see section above) or demographic characteristics (74). For example, older individuals put greater emphasis on positive feelings (75), which could translate to greater weighting of affect in seeking information. It is possible that what is a net benefit of disclosure for some individuals and groups will be a net cost for others. Information may be targeted and/or framed according to population characteristics such as age, education and mental health. Though much further research is needed in this domain, targeted disclosure could have far higher net welfare benefits than untargeted disclosure (76).

## Conclusion

It is increasingly possible for people to obtain information that bears on their future prospects, in terms of health, finance, and even romance. It is also increasingly possible for them to obtain information about the past, the present, and the future, whether or not that information bears on their personal lives. In principle, people's decisions about whether to seek or avoid information should depend on some integration of instrumental value, hedonic value, and cognitive value. But various biases can lead to both insufficient and excessive information-seeking. Individual differences in information-seeking may reflect different levels of susceptibility to those biases as well as varying emphasis on instrumental, hedonic and cognitive utility. Such differences may also be diagnostic of mental health.

Whether positive or negative, the value of information bears directly on significant decisions of government agencies, which are often charged with the task of calculating the welfare effects of mandatory disclosure, and which have long struggled with the task. Our hope is that the integrative framework of information-seeking motives offered here will facilitate these goals, and promote future research in this important domain.



## Glossary

**Information:** Data previously not known with complete certainty.

**Information-Seeking:** The active pursuit of knowledge, for example by asking questions, reading and conducting online searches.

**Curiosity:** The feeling of wanting to know. While curiosity is related to information-seeking the two concepts are distinct. In particular, it is possible to be curious but to avoid information, or to seek information despite a lack of curiosity. For example, a person may be curious whether they have a pre-disposition for cancer, but decide not to pursue such information in order to avoid negative feelings. A person may also seek financial information in order to make better financial decisions despite not being curious about such information.

**Instrumental Utility of Information:** A measure quantifying the amount by which information will enable achieving an end goal.

**Hedonic Utility of Information:** A measure quantifying the amount of pleasure (or other positive feeling) information would induce minus the amount of pain (or other negative feeling) it would induce.

**Cognitive Utility of Information:** A measure quantifying the degree to which information would strengthen internal mental models.

**Mental Models:** Representation of concepts and the relationship among them, which are used to comprehend and anticipate reality.

**Affect:** A physiological reaction that varies in valence (positive/negative) and arousal.

### **Outstanding Questions & Future Directions**

1. Where and how are the three information-seeking utilities represented in the brain? Where and how are they integrated into a value that drives choice?
2. Which neurotransmitters modulate information-seeking? Do different neurotransmitters play a role in modulating the effect of the three different information-seeking utilities?
3. Are information-seeking choices at a group level best explained by models that include all three utilities or ones that include only one or two (after penalizing for number of parameters)? This can be tested in studies where the different utilities are manipulated on a trial-by-trial basis (e.g., 5,42), or in studies in which the utilities are measured using different scales.
4. Are individual differences in the relative impact of the three motives on information-seeking stable across domains of knowledge (e.g., health information, financial information etc.)? Are the relative weights affected both by “trait” and “state” (for example do they systematically alter under stress)?
5. How (in neurological, computation and psychological terms) do people estimate the impact of information on action, affect and cognition?
6. How is the value of information learned? Do classic learning models such as Bayesian models and Reinforcement Learning models apply?
7. Do the different motives for information influence different psycho-physiological mechanisms (such as attention or bodily states) in directing information-seeking?

8. What are the development trajectories of the three motives of information-seeking? Are certain motives more pronounced early/late in life than others?
9. Which motives of information-seeking are conserved across species?
10. How should policymakers assess the costs and benefits of mandatory disclosure policies, in such areas as health, safety, and the environment?

**Acknowledgments.** We thank Liron Rozenkrantz, Joe Marks, Filip Gesiarz, Chris Kelly, Isabella De Taranto, Caroline Charpentier and Yulin Wang for comments on previous versions of this manuscript. TS is supported by a Wellcome Trust Senior Research Fellowship 214268/Z/18/Z.

## References

1. Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. *Neuron*, 88(3), 449-460.
2. Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. *Psychological bulletin*, 116(1), 75.
3. Stigler G (1961) The economics of information. *J Polit Econ* 69:213–225.
4. Grant S, Kajii A, Polak B (1998) Intrinsic preference for information. *J Econ Theory* 83(2):233–259.
5. Charpentier, C. J., Bromberg-Martin, E. S., & Sharot, T. (2018). Valuation of knowledge and ignorance in mesolimbic reward circuitry. *Proceedings of the National Academy of Sciences*, 201800547.
6. Bromberg-martin ES, Hikosaka O (2009) Midbrain dopamine neurons signal preference for advance information about upcoming rewards. *Neuron* 63(1):119–126.
7. Eliaz K, Schotter A (2007) Experimental testing of intrinsic preferences for noninstrumental information. *Am Econ Rev* 97(2):166–169.
8. Berlyne DE (1957) Uncertainty and conflict: a point of contact between information-theory and behavior-theory concepts. *Psychol Rev* 1(6):329–339.
9. Kreps DM, Porteus EL (1978) Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46(1):185–200.
10. Ligneul, R., Mermillod, M., & Morisseau, T. (2017). From curiosity relief to epistemic surprise: complementary roles of the prefrontal cortex and the ventral striatum in the neural valuation of knowledge. *bioRxiv*, 157644.
11. Brydevall, M., Bennett, D., Murawski, C., & Bode, S. (2018). The neural encoding of information prediction errors during non-instrumental information seeking. *Scientific reports*, 8(1), 6134.
12. Kang MJ, et al. (2009) The wick in the candle of learning: Epistemic curiosity activates reward circuitry and enhances memory. *Psychol Sci* 20(8):963–973.
13. Smith D V., Rigney AE, Delgado MR (2016) Distinct reward properties are encoded via corticostriatal interactions. *Sci Rep* 6:20093.
14. Tricomi E, Fiez JA (2012) Information content and reward processing in the human striatum during performance of a declarative memory task. *Cogn Affect Behav Neurosci* 12(2):361–372.

15. Jessup RK, O'Doherty JP (2014) Distinguishing informational from value-related encoding of rewarding and punishing outcomes in the human brain. *Eur J Neurosci* 39(11):2014–2026.
16. Gruber MJ, Gelman BD, Ranganath C (2014) States of curiosity modulate hippocampus-dependent learning via the dopaminergic circuit. *Neuron* 84(2):486–496.
17. van Lieshout LLF, Vandenbroucke ARE, Müller NCJ, Cools R, de Lange FP (2018) Induction and relief of curiosity elicit parietal and frontal activity. *J Neurosci* 38(10):2816–17.
18. Bromberg-Martin ES, Hikosaka O (2011) Lateral habenula neurons signal errors in the prediction of reward information. *Nat Neurosci* 14(9):1209–1216.
19. Blanchard TC, Hayden BY, Bromberg-Martin ES (2015) Orbitofrontal cortex uses distinct codes for different choice attributes in decisions motivated by curiosity. *Neuron* 85(3):602–614.
20. Schultz, Wolfram, Peter Dayan, and P. Read Montague. A neural substrate of prediction and reward. *Science* 275.5306 (1997): 1593-1599.
21. Thornton RL (2008) The demand for, and impact of, learning HIV status. *Am Econ Rev* 95(5):1829–1863.
22. Persoskie A, Ferrer RA, Klein WMP (2014) Association of cancer worry and perceived risk with doctor avoidance: an analysis of information avoidance in a nationally representative US sample. *J Behav Med* 37:977–987.
23. Dwyer LA, Shepperd JA, Stock ML (2015) Predicting avoidance of skin damage feedback among college students. *Ann Behav Med* 49:685–695.
24. Caplin A, Leahy J (2001) Psychological expected utility theory and anticipatory feelings. *Q J Econ* 116:55–79.
25. Koszegi B (2010) Utility from anticipation and personal equilibrium. *Econ Theory* 44:415–444.
26. Golman R, Hagman D, Loewenstein G (2017) Information avoidance. *J Econ Lit* 55(1):96–135.
27. Brunnermeier, Markus K. and Parker, Jonathan A. (2005). Optimal expectations. *American Economic Review*, 95(4), 1092-1118.
28. Golman, Russell and Loewenstein, George F. and Molnar, Andras and Saccardo, Silvia. (2019). The Demand for, and Avoidance of, Information. Available at SSRN
29. Golman & Loewenstein (2018). Information Gaps: A Theory of Preferences Regarding the Presence and Absence of Information.
30. Hertwig, R., & Engel, C. (2016). Homo ignorans: Deliberately choosing not to know. *Perspectives on Psychological Science*, 11(3), 359-372.
31. Bénabou, Roland and Tirole, Jean. (2002). Self-confidence and personal motivation. *The Quarterly Journal of Economics*, 117(3), 871-915.
32. Karlsson N, Loewenstein G, Seppi D (2009) The ostrich effect: selective attention to information. *J Risk Uncertain* 38(2):95–115.
33. Lerman, C., C. Hughes, S. Lemon, D. Main, C. Snyder, C. Durham, S. Narod, H. Lynch. (1998) What You Don't Know Can Hurt You: Adverse Psychological Effects in Members of BRCA1-Linked and BRCA2-Linked Families Who Decline Genetic Testing, *Journal of Clinical Oncology*, XVI, 1650–1654.

34. Loewenstein, G. (1987). Anticipation and the valuation of delayed consumption. *The Economic Journal*, 97(387), 666-684.
35. Craik, K. J. W. (1967). *The nature of explanation* (Vol. 445). CUP Archive.
36. Johnson-Laird, P.N. (1983). *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Cambridge: Cambridge University Press.
37. Friston, K. (2010). The free-energy principle: a unified brain theory?. *Nature reviews neuroscience*, 11(2), 127.
38. Berlyne DE (1957) Uncertainty and conflict: a point of contact between information-theory and behavior-theory concepts. *Psychol Rev* 1(6):329–339.
39. Kreps DM, Porteus EL (1978) Temporal resolution of uncertainty and dynamic choice theory. *Econometrica* 46(1):185–200.
40. Nickerson, R. S. Confirmation bias: A ubiquitous phenomenon in many guises. *Rev. Gen. Psychol.* 2, 175 (1998).
41. Klayman, J. & Ha, Y.-W. Confirmation, disconfirmation, and information in hypothesis testing. *Psychol. Rev.* 94, 211 (1987).
42. Eliaz, K. and Schotter, A. (2010). Paying for confidence: An experimental study of the demand for non-instrumental information. *Games and Economic Behavior*, 70(2), 304-324.
43. Kobayashi, K., & Hsu, M. (2018). Common Neural Code for Reward and Information Value. *bioRxiv*, 324665.
44. Weinstein, N. D. (1980). Unrealistic optimism about future life events. *Journal of personality and social psychology*, 39(5), 806.
45. Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological review*, 80(4), 237.
46. Langer, E. J. (1975). The illusion of control. *Journal of personality and social psychology*, 32(2), 311.
47. Sharot, T. (2011). The optimism bias. *Current biology*, 21(23), R941-R945.
48. Gilbert, D. T., Driver-Linn, E., & Wilson, T. D. (2002). The trouble with Vronsky: Impact bias in the forecasting of future affective states.
49. Sloman, S., & Fernbach, P. (2018). *The knowledge illusion: Why we never think alone*. Penguin.
50. Pallier, G., Wilkinson, R., Danthiir, V., Kleitman, S., Knezevic, G., Stankov, L., & Roberts, R. D. (2002). The role of individual differences in the accuracy of confidence judgments. *The Journal of general psychology*, 129(3), 257-299.
51. Aderka, I. M., Haker, A., Marom, S., Hermesh, H., & Gilboa-Schechtman, E. (2013). Information-seeking bias in social anxiety disorder. *Journal of abnormal psychology*, 122(1), 7.
52. Hildebrand-Saints, L., & Weary, G. (1989). Depression and social information gathering. *Personality and Social Psychology Bulletin*, 15(2), 150-160.
53. Camp, C. J. (1986). I am curious-grey: Information seeking and depression across the adult lifespan. *Educational Gerontology*, 12(4), 375-384.
54. Locander, W. B., & Hermann, P. W. (1979). The effect of self-confidence and anxiety on information seeking in consumer risk reduction. *Journal of Marketing Research*, 268-274
55. Gray, K. M., & Tonge, B. J. (2001). Are there early features of autism in infants and preschool children?. *Journal of paediatrics and child health*, 37(3), 221-226.



56. Jensen, P. S., Mrazek, D., Knapp, P. K., Steinberg, L., Pfeffer, C., Schowalter, J., & Shapiro, T. (1997). Evolution and revolution in child psychiatry: ADHD as a disorder of adaptation. *Journal of the American Academy of Child & Adolescent Psychiatry*, 36(12), 1672-1681.
57. Holaway, R. M., Heimberg, R. G., and Coles, M. E. (2006). A comparison of intolerance of uncertainty in analogue obsessive-compulsive disorder and generalized anxiety disorder. *J. Anxiety Disord.* 20, 158–174.
58. Luking, K. R., Pagliaccio, D., Luby, J. L., & Barch, D. M. (2016). Reward processing and risk for depression across development. *Trends in cognitive sciences*, 20(6), 456-468.
59. McCabe, C., Cowen, P. J., & Harmer, C. J. (2009). Neural representation of reward in recovered depressed patients. *Psychopharmacology*, 205(4), 667-677.
60. Pizzagalli DA, Iosifescu D, Hallett LA, Ratner KG, Fava M: Reduced hedonic capacity in major depressive disorder: evidence from a probabilistic reward task. *J Psychiatr*
61. Alloy, L. B., & Abramson, L. Y. (1979). Judgment of contingency in depressed and nondepressed students: Sadder but wiser?. *Journal of experimental psychology: General*, 108(4), 441.
62. Nestler, E. J., & Carlezon Jr, W. A. (2006). The mesolimbic dopamine reward circuit in depression. *Biological psychiatry*, 59(12), 1151-1159.
63. Davie, C. A. (2008). A review of Parkinson's disease. *British medical bulletin*, 86(1), 109-127.
64. Grace, A. A. (2016). Dysregulation of the dopamine system in the pathophysiology of schizophrenia and depression. *Nature Reviews Neuroscience*, 17(8), 524.
65. Giancardo, L., Sanchez-Ferro, A., Arroyo-Gallego, T., Butterworth, I., Mendoza, C. S., Montero, P., ... & Estépar, R. S. J. (2016). Computer keyboard interaction as an indicator of early Parkinson's disease. *Scientific reports*, 6, 34468.
66. Shukla, P., & Solanki, R. (2013). Web based keystroke dynamics application for identifying emotional state. Institute of Engineering and Technology, DAVV, Indore, India in Nov.
67. Sunstein, C. R. (2018). *The cost-benefit revolution*. MIT Press.
68. Sunstein, C. R. (2019). Ruining popcorn? The welfare effects of information. *Journal of Risk and Uncertainty*, 1-22.
69. Food and Drug Administration, HHS. (2011). Required warnings for cigarette packages and advertisements. Final rule. *Federal register*, 76(120), 36628.
70. Food and Drug Administration, HHS. (2014). Food labeling; nutrition labeling of standard menu items in restaurants and similar retail food establishments. Final rule. *Federal register*, 79(230), 71155.
71. Thunström, L. (2019). Welfare effects of nudges: The emotional tax of calorie menu labeling. *Judgment and Decision Making*, 14(1), 11.
72. Woolley, K. and Risen, J. L. (2018). Closing your eyes to follow your heart: Avoiding information to protect a strong intuitive preference. *Journal of personality and social psychology*, 114(2), 230.
73. Schneider B. (2017) *The Evolution of Airline Safety Videos*. City Lab <https://www.citylab.com/life/2017/12/the-evolution-of-airline-safety-videos/548858/>
74. Sakaki, M., Yagi, A., & Murayama, K. (2018). Curiosity in old age: A possible key to achieving adaptive aging. *Neuroscience & Biobehavioral Reviews*.

75. Carstensen, L. L. (1992). Social and emotional patterns in adulthood: support for socioemotional selectivity theory. *Psychology and aging*, 7(3), 331.
76. Allcott, H., & Kessler, J. B. (2019). The welfare effects of nudges: A case study of energy use social comparisons. *American Economic Journal: Applied Economics*, 11(1), 236-76.