Information-Seeking and Well-Being

Christopher A. Kelly

A dissertation submitted in partial fulfilment
of the requirements for the degree of

Doctor of Philosophy
of
University College London

Department of Experimental Psychology
University College London

October 10th, 2023

Declaration

I, Christopher Kelly, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the thesis.

Abstract

In today's '*information age'*, understanding how information consumption impacts well-being is vital. This thesis investigates the link between information-seeking behaviour and well-being, blending theory with empirical evidence.

Chapter 2 tests Sharot and Sunstein's (2020) information-seeking theory. Across five studies, I demonstrate that information-seeking is driven by Instrumental, Hedonic, and Cognitive Utility. The weights individuals assign to these different motives when searching for information remain relatively stable over time and correlate with mental health.

Chapter 3 explores the relationship between online information-seeking and well-being. Through four studies, I find that the emotional tone of online content shapes and reflects well-being. The emotional tone of websites users browse was analysed using natural language processing. It was found that negative content exposure was associated with decreased well-being. Experiments in which the valence of web content was manipulated and assessed mood, established a bi-directional causal relationship between the two.

Chapter 4 investigates how public and private stressors affect online search behaviours across three studies. Under stress, web queries shifted towards information than can guide action. For instance, during the pandemic, *"How"* queries on Google markedly increased. The number of *"How"* searches correlated with stress levels reported by 17K individuals weekly. An additional study showed that personal stressful events also selectively increased *"How"* searches. These patterns may be used as potential indicators of stress levels in a population.

Chapter 5 introduces a Google Chrome plugin that scores webpages on Instrumental, Hedonic, and Cognitive Utility. This tool aids users in refining their online journey, potentially reducing the web's adverse effects on well-being and enhancing user experience.

In conclusion, this thesis underscores the intricate tie between information-seeking and well-being in today's digital age. Understanding search motivations, acknowledging the emotional impact of content, and utilising tools for mindful choices can pave the way for healthier digital navigation.

Impact Statement

The quest for information is a fundamental part of the human existence, from ancient societies studying nature's patterns to the present-day '*information age*' driven by technological advancements. In this era of unprecedented data availability, individuals often face decisions about what information to seek or avoid. This thesis examines the motivations behind information-seeking choices and the consequences for individual well-being.

First, this thesis tests an integrative theory of information-seeking. This theoretical framework provides a comprehensive understanding of how people determine which information to seek and which to avoid. It suggests that people take into account the Instrumental, Hedonic, and Cognitive utilities of information when deciding what information to seek out. The findings can inform policy makers in tailoring campaigns to individuals' information preferences.

Importantly, this research identifies a bi-directional relationship between online information-seeking patterns and well-being. It shows that the types of information individuals choose to engage with can both shape and reflect their well-being. These insights can inform the development of tools for early detection of mental health problems from analysis on web browsing behaviour.

Extending from this foundational understanding, the research ventured into the practical wolrd of online navigation. Recognising the influence of digital content on mental health, I introduced a browser plugin tool designed to enhance informed online decision-making. This tool scores search engine results based on the Instrumental, Hedonic, and Cognitive utilities of the content on a webpage. It aims to enable users to make decisions that align with their browsing objectives and to mitigate the exposure to unhealthy information, leading to an enhanced overall online experience.

In conclusion, this thesis highlights the intricacies of information-seeking behaviour, revealing both the motivations behind people's online decisions and the resulting implications for well-being. By bridging theoretical insights with tangible, practical solutions, the study not only increases our understanding of information-seeking in the digital age but also sets forth a proactive approach to enhancing online experiences.

Acknowledgments

First and foremost, I want to express my greatest appreciation to my PhD advisor, Professor Tali Sharot. Her guidance has been invaluable, and the support and opportunities she provided have deeply influenced my professional journey. It's under her mentorship that I've grown into the scientist I am today.

I'm also grateful to the Affective Brain Lab, both past and present members. Your consistent support and camaraderie over the last five years have been truly enriching. A special mention to Laura Globig for being an essential pillar of support throughout this journey.

To my friends in London (Chris G, Sarah, Bryan, Karl, and Nils) and in Boston (Will, Jitske, Johnny, Dani, Baptiste, and Ippo): your companionship has enhanced my experiences over the years, and I'm thankful for all the moments we've shared.

Lastly, my family's unwavering support has been the foundation of this endeavour. Barrie and Nicky, my brothers, have been a source of constant encouragement. Most importantly, my deepest appreciation goes to my parents who have always motivated me to chase my dreams.

Notes To Examiners

The findings from Chapter 2 have been published in a peer-reviewed article: Kelly, C.A, & Sharot, T. (2021). Individual differences in information-seeking. *Nature Communications*, 12(1), 1-13.

The findings from Chapter 3 & 4 are currently under revision at peer-reviewed journals and have also been posted to pre-print repositories online: Kelly, C. A., & Sharot, T. (2023). Knowledge-Seeking Reflects and Shapes Well-Being.; Kelly, C. A., Blain, B., & Sharot, T. (2022). High-Level Characteristics of Web Queries Change Under Threat.

Finally, the technology discussed in Chapter 5 has been submitted to the United States Patent Office through MIT's technology licensing office. Both Professor Tali Sharot and I are listed as co-inventors.

Table of Contents

Abstract **4**

Impact Statement **5**

Acknowledgments **6**

Notes To Examiners **7**

Table of Figures **10**

Table of Tables **11**

Chapter 1 .... **12**

Introduction **12**
- 1.1 Information-Seeking 12
- 1.2 Current Theories of Information-Seeking 12
- 1.3 An Integrative Theory of Information-Seeking 14
- 1.4 Information-Seeking and Well-Being. 16
- 1.5 Facilitating Web-Browsing Searches 20
- 1.6 Summary 21
- 1.7 References 22

Chapter 2 .... **28**

Individual Differences in Information-Seeking **28**
- 2.1 Overview 28
- 2.2 Methods 30
- 2.3 Results 38
- 2.4 Discussion 51
- 2.5 References 54

Chapter 3 .... **58**

Information-Seeking Reflects and Shapes Well-Being **58**
- 3.1 Overview 58
- 3.2 Methods 59
- 3.3 Results 65
- 3.4 Discussion 76
- 3.5 References 79

Chapter 4 .... **82**

High-Level Characteristics of Web Searches Change Under Stress **82**
- 4.1 Overview 82
- 4.2 Methods 84
- 4.3 Results 91
- 4.4 Discussion 100
- 4.5 References 103

Chapter 5 .... **107**

A Tool to Facilitate Web-Browsing **107**
- 5.1 Overview 107
- 5.2 Tool Development 109
- 5.3 Discussion 116
- 5.4 References 119

Chapter 6 .... **120**

General Discussion **120**
References 127

Supplementary Materials ***129***
Chapter 2 129
Chapter 3 152

Table of Figures

Figure 1.1. Integrative model of Information-Seeking 15

Figure 2.1. Information-Seeking Motives 39

Figure 2.2. Information-Seeking related to Psychopathology 41

Figure 2.3. Information-Seeking Motives, Study 2 45

Figure 2.4. Association between information-seeking and mental health is robust to replication 41

Figure 2.5. Information-Seeking Motives in the Financial Domain 47

Figure 2.6. Information-Seeking Motives in the Health Domain 49

Figure 3.1. Data collection and pre-processing pipeline 61

Figure 3.2. The valence of the whole text of webpages is reflective of the valence of the text participants attend to the most 67

Figure 3.3. Self-guided browsing of negative content online is associated with poorer mental health 69

Figure 3.4. Browsing more negatively valanced webpages is associated with worse mood before and after browsing 71

Figure 3.5. Bi-directional relationship between mood and the valence of information consumed 73

Figure 3.6. Novel online intervention decreases the amount of negative information browsed online 75

Figure 4.1. "How" questions are associated with guiding actions 92

Figure 4.2. High-level characteristics of web searches alter during the pandemic 94

Figure 4.3. Self-reported stress is selectively associated with an increase in "How" searches 96

Figure 4.4. Personal stressful events alter high-level characteristics of searches 100

Figure 5.1. Presentation of Scores 108

Figure 5.2. Process of Tool. 112

Figure 5.3. Performance Metrics of Actionability and Knowledge Models.. 115

Table of Tables

Table 4.1. Demographics of respondents in the UK COVID-19 Social Study.................87

Table 4.2. The total number of participants providing data during each calendar week in the UK COVID-19 Social Study...88

Table 4.3. Coding of COIVID-19 related confinement variables 89

Table 4.4. An exploratory analysis showed that there was no increase in the use of any other question words under stress relative to control. The use of "What" showed a decrease under stress relative to control..99

Table 5.1. Performance Metrics of the Logistic Regression Models Predicting Actionability and Knowledge labels (i.e., 0 = low Actionability/Knowledge; 1 = High Actionability/Knowledge) of webpages 113

Chapter 1

Introduction

1.1 Information-Seeking

Throughout human history, the quest for information has been central to our daily lives. We can trace this back to hunter-gatherer societies who gained invaluable information on foraging by studying the patterns of nature, discerning when and where nourishing plants would emerge (Lee & Daly, 1999). Fast forward to the present, and we find ourselves in the '*information age*' (Castells, 1996), where the advancement of technology grants us access to a staggering wealth of data at unprecedented ease. This includes personalised insights into one's health through genetic tests, financial profiles via credit scores, and glimpses into personal lives through platforms like social media, to name a few. Amidst this vast sea of information, individuals are constantly faced with choices about what information to seek and what to avoid.

Considering the widespread nature of information-seeking decisions in our daily lives, it becomes crucial to comprehend how individuals determine the information they decide to seek or avoid, and how these choices impact their overall well-being. By information-seeking, I refer to the proactive quest for knowledge, such as asking questions, reading, and searching online (Sharot & Sunstein, 2020). Despite the significance of this issue, we lack a comprehensive understanding of what motivates information-seeking choices and the reasons behind individuals' differing preferences. For instance, a recent study (Sunstein, 2019) demonstrated an approximately fifty-fifty split among respondents wanting to know or avoid knowing information about their potential predisposition to cancer, the projected global temperature in 2100, or the calorie content of meal options.

1.2 Current Theories of Information-Seeking

Some of the earliest theories of information-seeking came from the field of Economics, which proposed that individuals seek information that offers rewards or prevents harm, or that has Instrumental Utility (Stigler, 1961; Hirshleifer & Ryley, 1979). To illustrate, consider a scenario where someone discovers that they have high cholesterol, a condition associated with an increased risk of heart-related issues. With this knowledge, they can make an informed decision to take preventive medication, thereby safeguarding their health. Thus, such information has high Instrumental Utility. More recent research has provided further evidence in support of individuals' desire for instrumental information (Kobayashi & Hsu, 2019; Wilson et al., 2014; Golman et al., 2021; Cogliati Dezza et al., 2022).

However, it's clear that Instrumental Utility is not the only driver behind information-seeking decisions, since individuals frequently seek information that doesn't necessarily change outcomes (Grant et al., 1998; Charpentier et al., 2018; Bromberg-Martin & Hikosaka, 2009). For example, individuals might want information about the life of Tutankhamun or the Irish Famine 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 (Grant et al., 1998). Having this perspective might be advantageous, as even if information seems irrelevant now, it could prove beneficial in the future (Berlyne, 1957; Kreps & Porteus, 1978).

Contradicting the notion that information is always beneficial, there are instances where individuals opt to avoid information with clear Instrumental Utility, such as medical diagnoses (Caplin & Eliaz, 2003; Lerman et al., 1998). For example, some individuals who are at risk of HIV, the virus that causes AIDS, avoid being tested, even when those tests have no financial cost (Caplin & Eliaz, 2003, Thornton, 2008; Persoskie et al., 2014; Dwyer et al., 2015).

This aversion to certain types of information has led researchers to investigate the Hedonic Utility of information, that is, how information can impact our emotional state (Persoskie et al., 2014; Pictet et al., 2011). Studies have documented that information can elicit both positive and negative emotions (Persoskie et al., 2014; Pictet et al., 2011). A classic example is learning about a genetic predisposition to certain cancers, which can evoke feelings of sadness and fear (Persoskie et al., 2014).

Given this observation, it's logical that individuals tend to seek information more when they anticipate positive outcomes rather than negative ones. This tendency aligns with many research findings (Sharot & Sunstein, 2020; Stigler, 1961; Hertwig & Engel, 2016; Persoskie et al., 2014; Golman et al., 2017; Karlsson et al., 2009; Lerman et al., 1998; Kobayashi et al., 2019; Charpentier et al., 2018). For instance, investors are more likely to monitor their portfolios when they predict an increase in value (Karlsson et al., 2009). Furthermore, studies have shown that people are more willing to pay for information signalling favourable news, such as financial gains (Charpentier et al., 2018). Conversely, people have been found to be willing to pay to avoid information when they anticipate bad news, such as information indicating financial loss (Charpentier et al., 2018). Ultimately, our emotional state plays a pivotal role in shaping our information-seeking behaviour and preferences.

Beyond the Instrumental and Hedonic Utility of information, individuals may seek information to enhance their understanding of the world (Sharot & Sunstein, 2020; Wilson et al., 2014; Oudeyer & Smith, 2016; Cogliati Dezza et al., 2017; Gershman, 2018; Schwartenbeck et al., 2019). For instance, Sharot & Sunstein (2020) have theorised that people will seek more information about concepts they think of often. This is because such information is especially relevant to their internal representation of their world and highly connected with many other concepts (Sharot & Sunstein, 2020). For example, someone who often thinks about dogs might be more intrigued about the relationship between dogs and wolves than someone who seldom thinks about dogs.

Another strategy that people may adopt in order to increase their comprehension of the world is to seek information about things which they are uncertain about (Wilson et al., 2014; Oudeyer & Smith, 2016; Cogliati Dezza et al., 2017; Gershman, 2018; Schwartenbeck et al., 2019; Afifi & Weiner, 2004; Berlyne, 1957; Kreps & Porteus, 1978; Nickerson, 1998; Klayman & Ha, 1987; Kappes et al., 2020; Stigler, 1961; Kobayashi et al., 2019; Charpentier et al., 2018; Hirshleifer & Riley, 1979; van Lieshout et al., 2018; Trudel et al., 2021). While some individuals seek information to confirm what they already believe (i.e., confirmation bias; Nickerson, 1998; Klayman & Ha, 1987; Kappes et al., 2020), others seek information when uncertain (Afifi & Weiner, 2004; Berlyne, 1957; Kreps & Porteus, 1978).

Trudel and colleagues (2021) suggest that the tendency to seek information, whether for confirmation or to reduce uncertainty, may vary depending on the environment. Together, this research highlights the complex relationship between certainty, uncertainty, and the pursuit of knowledge.

1.3 An Integrative Theory of Information-Seeking

In a recent theoretical paper, Sharot and Sunstein (2020) proposed an integrative theory of information-seeking combining the key empirical factors that drive information-seeking behaviour described above. Specifically, they hypothesised that when deciding whether to seek information, people first estimate what the information will reveal and then estimate the expected impact of that information on their Action, Affect and Cognition (**see Figure 1.1**). The estimated impact of information on Action, Affect and Cognition is referred to as Instrumental Utility, Hedonic Utility and Cognitive Utility, respectively (Sharot & Sunstein, 2020).

With regards to Action, the prediction is that people want information more when it can aid in selecting an action that will help gain rewards and avoid harm or is high in Instrumental Utility (Sharot & Sunstein, 2020; Stigler, 1961; Hirshleifer & Ryley, 1979; Kobayashi & Hsu, 2019; Wilson et al., 2014; Golman et al., 2021; Cogliati Dezza et al., 2022). For example, people would be more likely to want to know about automobile safety ratings if they are about to buy a car, as the information can inform their purchasing decision. With regards to Affect, all else being equal, people will be more likely to want information when they expect knowledge to make them feel better than ignorance (and vice versa, or is high in Hedonic Utility; Sharot & Sunstein, 2020; Stigler, 1961; Hertwig & Engel, 2016; Persoskie et al., 2014; Golman et al., 2017; Karlsson et al., 2009; Lerman et al., 1998; Kobayashi et al., 2019; Charpentier et al., 2018). For example, the prediction is that a student would be more likely to want to know their mark on an exam if they believe they have done well. With regards to Cognition, I propose that people will want information about concepts they think of often (Sharot & Sunstein, 2020). This is because such information is especially relevant to their internal representation of their world and highly connected with many other concepts (Sharot & Sunstein, 2020). For example, the prediction is that a person who thinks about dogs frequently, would be more interested in learning whether dogs are related to wolves compared to someone who rarely thinks about dogs.

It is also possible that people will seek information to reduce their uncertainty (Afifi & Weiner, 2004; Berlyne, 1957; Kreps & Porteus, 1978) or, conversely, to confirm their beliefs (Nickerson, 1998; Klayman & Ha, 1987; Kappes et al., 2020). Therefore, this factor should also be considered.

The estimates for Instrumental, Hedonic and Cognitive Utility of information can be positive (increasing information seeking), negative (increasing information avoidance) or zero (inducing indifference; Sharot and Sunstein, 2020). The hypothesis is that these estimates are integrated into a computation of the value of information, which will trigger information seeking or its active avoidance (Sharot & Sunstein, 2020). In Chapter 2, I will test this theory overtime and across domains—testing when people want information about personal-traits, finance and health.

Sharot & Sunstein (2020) further proposed that each of the three factors may be weighted differently, influencing the decision to seek or avoid information to different degrees (**see Figure 1.1**). Individual differences in information seeking may be related to the different weight individuals assign to each motive. For example, certain individuals may care most about the Instrumental Utility of information, whereas others may care most about the need to regulate their affective state, while others may assign equal weight to all three motives when seeking information, etc. In Chapter 2, I will quantify those differences and examine to what degree they are stable, or change, over time within and across domains, by conducting three longitudinal studies.

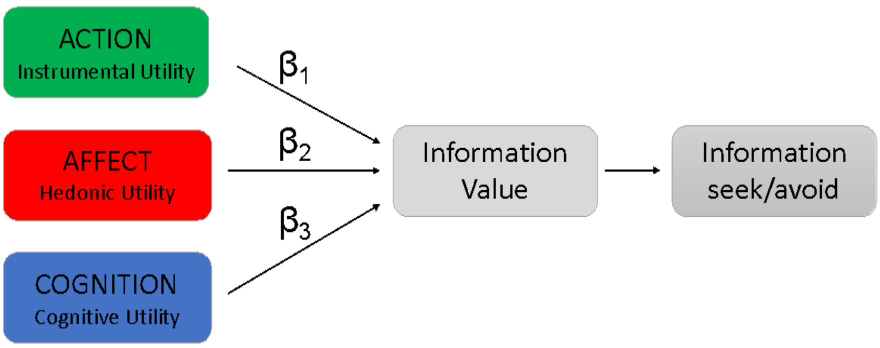

Figure 1.1. Integrative model of Information-Seeking. Information seeking and its avoidance is hypothesised to be driven by Instrumental Utility, Hedonic Utility and Cognitive Utility (Sharot & Sunstein, 2020). These values reflect the predicted impact of information on action, affect and cognition, respectively. These estimates are hypothesised to be integrated into a computation of the value of information, with different weights ($\beta_{1\text{–}3}$) assigned to each of the three factors. The integrated value can lead to information seeking or avoidance.

1.4 Information-Seeking and Well-Being.

Extant research investigating the association between information-seeking and mental health have primarily assessed the frequency of information-seeking, an approach which has led to mixed results (Aderka et al., 2013; Hildebrand-Saints & Weary, 1989; Camp, 1986; Locander & Hermann, 1979). For example, Aderka and colleagues (2013) examined the role of excessive reassurance-seeking in anxiety disorders, finding that frequent information-seeking can perpetuate symptoms of anxiety. In contrast, Hildebrand-Saints & Weary (1989) observed an adaptive role of information-seeking in managing stress, suggesting that individuals who actively seek information may better cope with stressors.

Here, we examined for an association between mental health and participants' motives for seeking information. Moreover, instead of using traditional psychopathology nosology, we adopted a dimensionality approach (Cuthbert & Insel, 2010; Cuthbert & Insel, 2013; Gillan et al., 2016). This approach deviates from the traditional classification of mental health disorders, such as those in the Diagnostic and Statistical Manual of Mental Disorders (DSM), which create distinct categories such as depressive and anxiety disorders (APA, 2013). This approach, although being useful in standardising diagnosis, has limitations such as the comorbidity and heterogeneity of diagnoses. The transdiagnostic approach, instead, focusses on these shared mechanisms allowing it to offer a more nuanced understanding to the etiology of mental health. For example, Gillan and colleagues (2016) found that many disorders share common cognitive deficits, such as impaired decision-making (e.g., OCD, eating disorder, impulsivity. By applying a transdiagnostic approach, via inputting individual items from many psychopathology questionnaires into a factor analysis, they identified a distinct factor that explained these deficits, which largely contained symptoms related to all these disorders. In addition, Caspi and colleagues (2014) have proposed the concept of the p-factor, which represents a general dimension of psychopathology. Their research suggests that this factor underlies a wide range of mental health disorders.

If indeed there are individual differences in the importance people place on different types of information, such differences might be related to well-being (Sharot & Sunstein, 2020). Many psychopathology symptoms can be broadly characterized as problems in affective processes, cognitive functions, and action planning and execution (Cuthbert & Insel, 2010; Cuthbert & Insel, 2013). Thus, abnormalities in these domains may reveal themselves in the type of information people choose to seek or avoid. For instance, poor well-being is characterised by a reduction in the belief that one has agency over outcomes (Alloy & Abramson, 1979; Yoshie & Haggard, 2013), which may lead to a reduction in the impact of instrumental utility on information-seeking. Additionally, people with poor well-being have been shown to have a negativity bias in the context of information-seeking (Owens et al., 2004; Aderka, et al., 2013), that is they tend to attend to negative information more than those with greater well-being. Thus in the context of the integrative information-seeking framework, they may be more likely to seek information that they attribute a negative hedonic value to.

Moreover, I hypothesise that the relationship between information-seeking and well-being is bidirectional. For example, negative thoughts may lead to searches for information with a similar sentiment, resulting in the consumption of negatively valenced content, which could in turn exacerbate one's negative affective state. This potential mechanism is consistent with findings suggesting that people with depression tend to engage with stimuli that perpetuate their sadness (Milgram et al., 2015), and is analogous to the mechanism hypothesised to underlie rumination. Specifically, it has been suggested that continuous negative thoughts (akin to internal information-seeking) can sustain and exacerbate low moods through a feedback loop (Watkins, 2008; Michl et al., 2013). Empirically testing this hypothesis is especially important today, given the exponential increase in the availability, speed, and ease of access to information, which likely amplifies the impact of information-seeking patterns on mental health.

The relationship between information-seeking and well-being may be context dependent. In particular, different contexts may trigger different information-seeking patterns, which in turn, may have implications on well-being. Take, for instance, the transformative global event of the Covid-19 outbreak. Such a seismic shift in the environment may prompt individuals to seek specific types of information (Charpentier et al., 2022). In Chapters 2, 3 and 4, I will examine the relationship between information-seeking and well-being within three domains: (i) when seeking self-referential information (i.e., personal traits), (ii) during web-browsing, and (iii) under stress.

First, in Chapter 2, I will test whether the weights people assign to Action, Affect and Cognition when faced with self-referential information are related to self-reported mental health (Sharot & Sunstein, 2020). The rationale for testing the relationship between information-seeking and well-being specifically in the domain of self-referential information is because poor mental health is often associated with problems related to self-perception and thoughts regarding the self (Hards et al., 2020; Christensen et al., 2003; Graham et al., 2014; Sass & Parnas, 2003; Jacobi et al., 2004; Doron et al., 2008; Silverstone & Salsali, 2003).

After testing the relationship between information-seeking and well-being in the domain of self-referential information, in Chapter 3, I will shift the lens to a broader landscape: the internet. The surge in online activity, especially in recent years, underscores the urgency to understand the relationship between web-browsing patterns and well-being (DataReportal, 2022). Despite its importance, surprisingly little is known about how online information-seeking relates to our well-being. Most research on online behaviour and well-being at an individual level has predominantly focused on assessing screentime (Babic et al., 2017; Page et al., 2010; Granic et al., 2014; Odgers, 2018). Interestingly, this line of research has yielded mixed findings. Some studies have shown a relationship between increased screentime and well-being (Sanders et al., 2024), while others have found no effect (Vanman et al., 2018; Allcott et al., 2019). Another line of research has focused on assessing the characteristics of what people *share* online (De Choudhury et al., 2013; Kelley & Gillan, 2022; Eichstaedt et al., 2018). This approach has led to more consistent findings, such as individuals with higher levels of depression posting more negatively-valenced tweets (Kelley & Gillan, 2022) and Facebook posts (Eichstaedt et al., 2018).

Complementary to previous information-sharing research, I will test the relationship between online information-seeking patterns and individual well-being. It is likely that these patterns will vary significantly among individuals, as observed in controlled lab studies (Kelly & Sharot, 2021; Kobayashi et al., 2019; Sunstein, 2019), and such variations may provide deep insights into one's well-being. I theorise that the affective properties of the information people choose to consume during self-guided searches reflect and shape their mental health, forming a feedback loop.

Regarding the first direction of this hypothesis—one's affective state influencing the affective characteristics of information sought online—this hypothesis aligns with findings that people with depression tend to engage with stimuli that perpetuate their sadness (Milgram et al., 2015). This mechanism is analogous to that hypothesised to underlie rumination (Watkins, 2008; Michl et al., 2013). In addition, this hypothesis is consistent with rich literature showing that affect alters information seeking and decision making (Sharot & Sunstein, 2020; Kelly & Sharot et al., 2021; Lerner et al., 2015; Hockey et al., 2000; George et al., 2016; Paulus & Angela, 2012; Phelps et al., 2014; Pictet et al., 2011; Stigler, 1961; Karlsson et al., 2009; Charpentier et al., 2018).

In the second direction—how the affective characteristics of browsed information impact well-being—previous research has shown a robust relationship between exposure to negative words and well-being (Mathews & MacLeod, 2005; Velte, 1968; Lyubomirsky et al., 1998). For example, Lyubomirsky and colleagues (1998) had participants read negative or positive passages and those who read negative passages reported lower mood levels afterward. Thus, it is reasonable to think that affective characteristics of words on webpages that people expose themselves to will have an impact on their well-being.

To quantify the affective properties of the information people expose themselves to online, we will ask individuals to share their web-browsing history and then use a natural language processing (NLP) approach to quantify the valence of the words on webpages that they browse. I will first relate these affective characteristics to participants' emotional and psychological well-being and then, if a significant relationship emerges, I will manipulate these factors to examine for a causal relationship. Finally, we will also examine whether providing cues about the potential emotional impact of webpages on well-being can influence participants' web-browsing behaviour, in a way that is consistent with improvements in well-being.

Thus, should a bi-directional relationship be established between the type of information that is consumed from self-guided web searches and mental health, it would have significant theoretical and practical implications. In particular, the digital nature of online activities simplifies assessment and opens up the potential for real-time practical applications. Knowledge of the relationship between online information-seeking patterns and mental health can inform the development of tools that could complement existing interventions, such as screen time awareness tools (Kim et al., 2016; Kovacs et al., 2022), and digital phenotyping methods (see Reece et al., 2017;

Otenen et al., 2023; Guntuku et al., 2020; Valdez et al., 2020; De Choudhury et al., 2013; Kelley & Gillan, 2022; Eichstaedt et al., 2018).

In Chapter 4, I explore how stressful life events impact information-seeking patterns. These stressful events may be global (e.g., war, pandemic) or unique to the individual (e.g., being diagnosed with cancer, losing one's job, divorce). Abundant research highlights that such events often lead to stress, anxiety, confusion, and a reduced sense of control, impacting mental and emotional well-being (Finlay-Jones & Brown, 1981; Francis et al., 2012; McLaughlin et al., 2010; Miloyan et al., 2018; Suls & Mullen, 1981; Globig et al., 2020). The American Psychological Association (APA; VandenBos, 2007) defines a stressful event as "*an occurrence or circumstance that individuals perceive as threatening, challenging, or demanding, thereby eliciting a stress response*." This broad definition encapsulates a range of experiences, from major life changes and daily hassles to situations such as job loss, financial difficulties, relationship conflicts, health issues, traumatic experiences, and environmental disasters (VandenBos, 2007).

One available adaptive reaction to stress is to seek information that can help guide action to promote adaptation (Hirshleifer & Riley, 1979; Sharot & Sunstein, 2020; Stigler, 1961). Such actions can be directly related to the event experienced (e.g., during wartime people may search for information on how to secure windows from being shuttered by rockets) or indirectly related (e.g., searching for activities that can distract oneself from the adversity). However, this strategy of seeking instrumental information may be specific when individuals have agency regarding the event causing stressed, as research indicates that having a sense of control can reduce stress (Bandura, 1997; Frazier et al., 2001). For instance, studies have shown that when individuals perceive they have control over a situation, they experience lower stress levels (Bandura, 1997; Frazier et al., 2001). However, in scenarios where individuals have little or no control over the outcome, such as in experiments where participants receive electric shocks, this strategy is less likely to be effective. In these cases, the lack of agency can exacerbate stress rather than alleviate it (Geer et al., 1970).

To date, research on the relationship between information-seeking and stress has mostly focused on the frequency of information-seeking. Some studies propose that stress is associated with greater information seeking (Drouin et al., 2020; Ebrahim et al., 2020; Loosen et al., 2021), which may decrease the sense of uncertainty that is heightened under stress. Others, however, suggest that stress leads to an avoidance reaction which is characterised by less information-seeking about the stressor (Kash et al. 2000; Chae, 2016).

I take a different theoretical viewpoint. Rather than focusing on whether stress generally enhances or reduces information seeking, I test the hypothesis that when experiencing abrupt stressful life events people are more likely to search for information that can direct action – a reaction which may be adaptive. In other words, stress may alter the type, rather than frequency, of information people seek. Across multiple studies, I test whether such changes can be detected and quantified using NLP analysis of web-browsing searches obtained from both population level Google Trends data as well as through controlled web-browsing study assessing individuals web-browsing behaviour. I propose that by examining the features of the information that people seek, we can gain insight into their external state and their internal reaction to that state.

To that end, I also examine whether the negative context people find themselves in is associated with a change in the valence of web searches. That is, whether a negative state may lead to more negatively valenced searches due to that state, or alternatively to more positive searches in an attempt perhaps to counter the negative state. The former possibility is supported by studies showing that anxious individuals have a bias towards negative stimuli (Bar-Haim et al., 2007; Cisler et al., 2010; MacLeod et al., 1988), which suggests that anxiety may increase the search for negative information.

1.5 Facilitating Web-Browsing Searches

Search results are shaped by opaque algorithms that do not necessarily align with users' goals (Rainie, Lee & Anderson, 2017). Consequently, individuals dedicate countless hours to absorbing information that may not yield practical benefits, and in some cases, may have a detrimental effect on their well-being (Kelly & Sharot, 2023). For example, by consuming negatively valenced information that is not informative or helpful.

In Chapter 5, informed by the insights from Chapters 2-4, I introduce a tool designed to guide individuals in their online information-seeking, with the intent to enhance decision-making, improve mental health, and enrich understanding. Analogous to how nutritional labels inform us about the contents of our food, such as calorie and fat content, our tool offers *'content labels'* for webpages listed in search engine results, allowing users to discern the characteristics of information before engaging with it.

In particular, the software (in the form of a Google Chrome plugin) informs users of three properties that can guide information-consumption decisions: (i) actionability (the ability of text on a webpage to guide action, on average); (ii) ability of text on a webpage to enhance understanding, on average; (iii) sentiment (e.g., how positive or negative the text on a webpage is). These three properties were selected based on empirical research that indicates that people's key motives for seeking information is to (i) guide their actions and decisions (Kelly & Sharot, 2021; Cogliati Dezza et al., 2022; Stigler, 1961; for review Sharot & Sunstein, 2020), (ii) improve comprehension (Cogliati Dezza et al., 2022; Sharot & Sunstein, 2020) and (iii) improve affect (Kelly & Sharot, 2021; Sharot & Sunstein, 2020; Charpentier et al., 2018; Loewenstein, 1994; Caplin & Leahy, 2001; Kőszegi, 2010; Golman et al., 2017).

1.6 Summary

The pursuit of information has been integral to human existence since ancient times, evolving from the knowledge-gathering practices of early societies to the present-day '*information age'*. Today, technology has granted unprecedented access to an overwhelming volume of data, including personalised health insights and financial statuses. In this vast sea of information, individuals are continually confronted with decisions on what to seek and what to avoid.

Given the ubiquity of such information-seeking choices, it's imperative to understand the motivations behind these decisions and assess their implications for individual well-being. Information-seeking is an active behaviour involving various approaches like reading newspapers, and online browsing. Yet, there remains a gap in our understanding of the underlying factors that guide people's choices.

We examine an integrative theory of information-seeking that combines key factors—Instrumental Utility, Hedonic Utility, Cognitive Utility (Sharot & Sunstein, 2020). This theory proposes that individuals estimate these three values before making a choice to either seek or avoid information. In the context of Instrumental Utility, the hypothesis is that people will be inclined to acquire information when it can lead to reward or the avoidance of harm. For instance, someone contemplating the purchase of a home would want data on the neighbourhood's crime rate, leveraging this information for an informed transaction. Regarding Hedonic Utility, the theory suggests that people will seek information when it is perceived as more emotionally beneficial than remaining uninformed. For example, someone expecting a positive outcome from an interview will be more inclined to seek information regarding the outcome. Lastly, for Cognitive Utility, the theory posits that individuals seek information related to topics that are frequently thought about. This is because such information is especially relevant to their internal representation of their world and highly connected with many other concepts (Sharot & Sunstein, 2020). For example, someone often thinking about football might be more intrigued about the outcome of the weeks football games compared to someone who seldom thinks about football.

The theory also proposes that differences in information-seeking patterns will shape and reflect well-being. This means that the types of information people choose to consume can both reflect their current state of well-being and influence it.

With the aim of facilitating online decision-making, we introduce a browser plugin tool. This tool labels search engine results according to their actionability, capacity to increase comprehension about a topic, and emotional valence, thus potentially facilitating a more informed user engagement.

In summation, this thesis will examine the key motivations for seeking information and how different information-seeking patterns reflect and shape well-being. It holds the potential for enhanced mental health diagnostics and fosters a more mindful navigation in the '*information age'*.

1.7 References

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.

Afifi, W. A., & Weiner, J. L. (2004). Toward a theory of motivated information management. Communication theory, 14(2), 167-190.

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.

American Psychiatric Association, D. S. M. T. F., & American Psychiatric Association. (2013). Diagnostic and statistical manual of mental disorders: DSM-5 (Vol. 5, No. 5). Washington, DC: American psychiatric association.

Bar-Haim, Y., Lamy, D., Pergamin, L., Bakermans-Kranenburg, M. J., & Van Ijzendoorn, M. H. (2007). Threat-related attentional bias in anxious and non-anxious individuals: a meta-analytic study. Psychological bulletin, 133(1), 1.

Berlyne, D. E. (1957). Uncertainty and conflict: a point of contact between information-theory and behavior-theory concepts. Psychological Review, 64(6p1), 329.

Bromberg-Martin, E. S., & Hikosaka, O. (2009). Midbrain dopamine neurons signal preference for advance information about upcoming rewards. Neuron, 63(1), 119-126.

Camp, C. J. (1986). I am curious-grey: Information seeking and depression across the adult lifespan. Educational Gerontology, 12(4), 375-384.

Caplin, A., & Eliaz, K. (2003). Aids policy and psychology: A mechanism-design approach. RAND Journal of Economics, 631-646.

Caplin, A., & Leahy, J. (2001). Psychological expected utility theory and anticipatory feelings. The Quarterly Journal of Economics, 116(1), 55-79.

Castells, M. (1996). The information age: Economy, society and culture (3 volumes). Blackwell, Oxford, 1997, 1998.

Chae, J. (2016). Who avoids cancer information? Examining a psychological process leading to cancer information avoidance. Journal of health communication, 21(7), 837-844.

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, 115(31), E7255-E7264.

Christensen, P. N., Stein, M. B., & Means-Christensen, A. (2003). Social anxiety and interpersonal perception: A social relations model analysis. Behaviour Research and Therapy, 41(11), 1355-1371.

Cisler, J. M., & Koster, E. H. (2010). Mechanisms of attentional biases towards threat in anxiety disorders: An integrative review. Clinical psychology review, 30(2), 203-216.

Cogliati Dezza, I., Yu, A. J., Cleeremans, A., & Alexander, W. (2017). Learning the value of information and reward over time when solving exploration-exploitation problems. Scientific Reports, 7(1), 1-13.

Cuthbert, B. N., & Insel, T. R. (2010). Toward new approaches to psychotic disorders: the NIMH Research Domain Criteria project.

Cuthbert, B. N., & Insel, T. R. (2013). Toward the future of psychiatric diagnosis: the seven pillars of RdoC. BMC medicine, 11(1), 1-8.

DataReportal (2022), “Digital 2022 Global Digital Overview,” retrieved from https://datareportal.com/reports/digital-2022-global-overview-report
De Choudhury, M., Gamon, M., Counts, S. & Horvitz, E. (2013). Predicting depression via social media. ICWSM 2, 128–137.
Dezza, I. C., Maher, C., & Sharot, T. (2022). People adaptively use information to improve their internal states and external outcomes. Cognition, 228, 105224.
Doron, G., Moulding, R., Kyrios, M., & Nedeljkovic, M. (2008). Sensitivity of self-beliefs in obsessive compulsive disorder. Depression and Anxiety, 25(10), 874-884.
Drouin, M., McDaniel, B. T., Pater, J., & Toscos, T. (2020). How parents and their children used social media and technology at the beginning of the COVID-19 pandemic and associations with anxiety. Cyberpsychology, Behavior, and Social Networking, 23(11), 727-736.
Dwyer LA, Shepperd JA, Stock ML (2015) Predicting avoidance of skin damage feedback among college students. Ann Behav Med 49:685–695.
Ebrahim, A. H., Saif, Z. Q., Buheji, M., AlBasri, N., Al-Husaini, F. A., & Jahrami, H. (2020). COVID-19 information-seeking behavior and anxiety symptoms among parents. OSP Journal of Health Care and Medicine, 1(1), 1-9.
Eichstaedt, J. C., Smith, R. J., Merchant, R. M., Ungar, L. H., Crutchley, P., Preoţiuc-Pietro, D., ... & Schwartz, H. A. (2018). Facebook language predicts depression in medical records. Proceedings of the National Academy of Sciences, 115(44), 11203-11208.
Finlay-Jones, R., & Brown, G. W. (1981). Types of Stressful Life Event and the Onset of Anxiety and Depressive Disorders. Psychological Medicine, 11(4), 803–815. https://doi.org/10.1017/S0033291700041301
Francis, J. L., Moitra, E., Dyck, I., & Keller, M. B. (2012). The impact of stressful life events on relapse of generalized anxiety disorder. Depression and Anxiety, 29(5), 386–391. https://doi.org/10.1002/da.20919
Gershman, S. J. (2018). Deconstructing the human algorithms for exploration. Cognition, 173, 34-42.
Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A., & Daw, N. D. (2016). Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. Elife, 5, e11305.
Globig, L. K., Blain, B., & Sharot, T. (2022). Perceptions of personal and public risk: Dissociable effects on behavior and well-being. Journal of Risk and Uncertainty, 1-22.
Golman, R., Gurney, N., & Loewenstein, G. (2021). Information gaps for risk and ambiguity. Psychological Review, 128(1), 86.
Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information avoidance. In Journal of Economic Literature. https://doi.org/10.1257/jel.20151245
Graham, K. T., Martin-Iverson, M. T., Holmes, N. P., Jablensky, A., & Waters, F. (2014). Deficits in agency in schizophrenia, and additional deficits in body image, body schema, and internal timing, in passivity symptoms. Frontiers in Psychiatry, 5, 126.
Grant S, Kajii A, Polak B (1998) Intrinsic preference for information. J Econ Theory 83(2):233–259.
Hards, E., Ellis, J., Fisk, J., & Reynolds, S. (2020). Negative view of the self and symptoms of depression in adolescents. Journal of affective disorders, 262, 143-148.
Hertwig, R., & Engel, C. (2016). Homo ignorans: Deliberately choosing not to know. Perspectives on Psychological Science, 11(3), 359-372.

Hildebrand-Saints, L., & Weary, G. (1989). Depression and social information gathering. Personality and Social Psychology Bulletin, 15(2), 150-160.
Hirshleifer, J., & Riley, J. G. (1979). The analytics of uncertainty and information-an expository survey. Journal of economic literature, 17(4), 1375-1421.
Jacobi, C., Paul, T., de Zwaan, M., Nutzinger, D. O., & Dahme, B. (2004). Specificity of self-concept disturbances in eating disorders. International Journal of Eating Disorders, 35(2), 204-210.
Kappes, A., Harvey, A. H., Lohrenz, T., Montague, P. R., & Sharot, T. (2020). Confirmation bias in the utilization of others' opinion strength. Nature neuroscience, 23(1), 130-137.
Karlsson, N., Loewenstein, G., & Seppi, D. (2009). The ostrich effect: Selective attention to information. Journal of Risk and uncertainty, 38(2), 95-115.
Kash, K. M., Ortega-Verdejo, K., Dabney, M. K., Holland, J. C., Miller, D. G., & Osborne, M. P. (2000, June). Psychosocial aspects of cancer genetics: women at high risk for breast and ovarian cancer. In Seminars in Surgical Oncology (Vol. 18, No. 4, pp. 333-338). New York: John Wiley & Sons, Inc.
Kelley, S. W., & Gillan, C. M. (2022). Using language in social media posts to study the network dynamics of depression longitudinally. Nature communications, 13(1), 1-11.
Kelly, C. A., & Sharot, T. (2023). Knowledge-Seeking Reflects and Shapes Well-Being.
Kelly, C.A., Sharot, (2021) T. Individual differences in information-seeking. Nat Commun 12, 7062. https://doi.org/10.1038/s41467-021-27046-5
Klayman, J., & Ha, Y. W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. Psychological review, 94(2), 211.
Kobayashi, K., & Hsu, M. (2019). Common neural code for reward and information value. Proceedings of the National Academy of Sciences, 116(26), 13061-13066.
Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. Nature human behaviour, 3(6), 587-595.
Kőszegi, B. (2010). Utility from anticipation and personal equilibrium. Economic Theory, 44, 415-444.
Kreps, D. M., & Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. Econometrica: journal of the Econometric Society, 185-200.
Lee, R. B., Daly, R. H., & Daly, R. (Eds.). (1999). The Cambridge encyclopedia of hunters and gatherers. Cambridge University Press.
Lerman, C., Hughes, C., Lemon, S. J., Main, D., Snyder, C., Durham, C., ... & Lynch, H. T. (1998). What you don't know can hurt you: adverse psychologic effects in members of BRCA1-linked and BRCA2-linked families who decline genetic testing. Journal of Clinical Oncology, 16(5), 1650-1654.
Locander, W. B., & Hermann, P. W. (1979). The effect of self-confidence and anxiety on information seeking in consumer risk reduction.
Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. Psychological bulletin, 116(1), 75.
Loosen, A. M., Skvortsova, V., & Hauser, T. U. (2021). Obsessive–compulsive symptoms and information seeking during the Covid-19 pandemic. Translational psychiatry, 11(1), 309.
MacLeod, C., & Mathews, A. (1988). Anxiety and the allocation of attention to threat. The Quarterly journal of experimental psychology, 40(4), 653-670.
McLaughlin, K. A., Conron, K. J., Koenen, K. C., & Gilman, S. E. (2010). Childhood adversity, adult stressful life events, and risk of past-year psychiatric disorder: A test

of the stress sensitization hypothesis in a population-based sample of adults. Psychological Medicine, 40(10), 1647–1658.
Miloyan, B., Joseph Bienvenu, O., Brilot, B., & Eaton, W. W. (2018). Adverse life events and the onset of anxiety disorders. Psychiatry Research, 259, 488–492. https://doi.org/10.1016/j.psychres.2017.11.027
Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. Review of general psychology, 2(2), 175-220.
Otenen, E., Seckin, O. C., Valdez, D., & Bollen, J. (2023). Quantifying the Digital Phenotype of Loneliness on Twitter. In Proceedings of the Annual Meeting of the Cognitive Science Society (Vol. 45, No. 45).
Oudeyer, P. Y., Lopes, M., Kidd, C., & Gottlieb, J. (2016). Curiosity and intrinsic motivation for autonomous machine learning. ERCIM News, 107, 34-35.
Persoskie, A., Ferrer, R. A., & Klein, W. M. (2014). Association of cancer worry and perceived risk with doctor avoidance: an analysis of information avoidance in a nationally representative US sample. Journal of behavioral medicine, 37(5), 977-987.
Pictet, A., Coughtrey, A. E., Mathews, A., & Holmes, E. A. (2011). Fishing for happiness: The effects of generating positive imagery on mood and behaviour. Behaviour research and therapy, 49(12), 885-891.
Phelps, E. A., Lempert, K. M., & Sokol-Hessner, P. (2014). Emotion and decision making: multiple modulatory neural circuits. Annual review of neuroscience, 37, 263-287.
Rainie, L., & Anderson, J. (2017). Code-dependent: Pros and cons of the algorithm age.
Reece, A. G., Reagan, A. J., Lix, K. L., Dodds, P. S., Danforth, C. M., & Langer, E. J. (2017). Forecasting the onset and course of mental illness with Twitter data. Scientific reports, 7(1), 13006.
Sass, L. A., & Parnas, J. (2003). Schizophrenia, consciousness, and the self. Schizophrenia bulletin, 29(3), 427-444.
Schwartenbeck, P., Passecker, J., Hauser, T. U., FitzGerald, T. H., Kronbichler, M., & Friston, K. J. (2019). Computational mechanisms of curiosity and goal-directed exploration. Elife, 8, 1-45.
Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. Nature Human Behaviour, 4(1), 14-19.
Silverstone, P. H., & Salsali, M. (2003). Low self-esteem and psychiatric patients: Part I–The relationship between low self-esteem and psychiatric diagnosis. Annals of general hospital psychiatry, 2(1), 1-9.
Stigler, G. J. (1961). The Economics of Information. P. Journal of Political Economy, 69.
Suls, J., & Mullen, B. (1981). Life Events, Perceived Control and Illness: The Role of Uncertainty. Journal of Human Stress, 7(2), 30–34.
Sunstein, C. R. (2019). Ruining popcorn? The welfare effects of information. Journal of Risk and Uncertainty, 58(2), 121-142.
Thornton, R. L. (2008). The demand for, and impact of, learning HIV status. American Economic Review, 98(5), 1829-1863.
Trudel, N., Scholl, J., Klein-Flügge, M. C., Fouragnan, E., Tankelevitch, L., Wittmann, M. K., & Rushworth, M. F. (2021). Polarity of uncertainty representation during exploration and exploitation in ventromedial prefrontal cortex. Nature Human Behaviour, 5(1), 83-98.
van Lieshout, L. L., Vandenbroucke, A. R., Müller, N. C., Cools, R., & de Lange, F. P. (2018). Induction and relief of curiosity elicit parietal and frontal activity. Journal of Neuroscience, 38(10), 2579-2588.

VandenBos, G. R. (2007). APA dictionary of psychology. American Psychological Association.
Watkins, E. R. (2008). Constructive and unconstructive repetitive thought. Psychological bulletin, 134(2), 163.
Wilson, R. C., Geana, A., White, J. M., Ludvig, E. A., & Cohen, J. D. (2014). Humans use directed and random exploration to solve the explore–exploit dilemma. Journal of Experimental Psychology: General, 143(6), 2074-2081.
Guntuku, S. C., Sherman, G., Stokes, D. C., Agarwal, A. K., Seltzer, E., Merchant, R. M., & Ungar, L. H. (2020). Tracking mental health and symptom mentions on Twitter during COVID-19. Journal of general internal medicine, 35, 2798-2800.
Valdez, D., Ten Thij, M., Bathina, K., Rutter, L. A., & Bollen, J. (2020). Social media insights into US mental health during the COVID-19 pandemic: Longitudinal analysis of Twitter data. Journal of medical Internet research, 22(12), e21418.
Kovacs, V. A., Starc, G., Brandes, M., Kaj, M., Blagus, R., Leskošek, B., ... & Okely, A. D. (2022). Physical activity, screen time and the COVID-19 school closures in Europe–An observational study in 10 countries. European journal of sport science, 22(7), 1094-1103.
Kim, Y. H., Jeon, J. H., Choe, E. K., Lee, B., Kim, K., & Seo, J. (2016). TimeAware: Leveraging framing effects to enhance personal productivity. In Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems (pp. 272-283).
Lyubomirsky, S., Caldwell, N. D., & Nolen-Hoeksema, S. (1998). Effects of ruminative and distracting responses to depressed mood on retrieval of autobiographical memories. Journal of personality and social psychology, 75(1), 166.
Velte, P., & Stiglbauer, M. (2012). Audit market concentration and its influence on audit quality. International Business Research 5.11 (2012): S. 146-161. 25.03. 2013.
Mathews, A., & MacLeod, C. (2005). Cognitive vulnerability to emotional disorders. Annu. Rev. Clin. Psychol., 1, 167-195.
Karlsson, N., Loewenstein, G., & Seppi, D. (2009). The ostrich effect: Selective attention to information. Journal of Risk and uncertainty, 38, 95-115.
Paulus, M. P., & Angela, J. Y. (2012). Emotion and decision-making: affect-driven belief systems in anxiety and depression. Trends in cognitive sciences, 16(9), 476-483.
George, J. M., & Dane, E. (2016). Affect, emotion, and decision making. Organizational Behavior and Human Decision Processes, 136, 47-55.
Hockey, G. R. J., John Maule, A., Clough, P. J., & Bdzola, L. (2000). Effects of negative mood states on risk in everyday decision making. Cognition & Emotion, 14(6), 823-855.
Lerner, J. S., Li, Y., Valdesolo, P., & Kassam, K. S. (2015). Emotion and decision making. Annual review of psychology, 66, 799-823.
Allcott, Hunt, Matthew Gentzkow, and Chuan Yu. 2019. "Trends in the diffusion of misinformation on social media." Research & Politics 6.2.
Vanman, E. J., Baker, R., & Tobin, S. J. (2018). The burden of online friends: The effects of giving up Facebook on stress and well-being. The Journal of social psychology, 158(4), 496-508.
Sanders, T., Parker, P. D., del Pozo-Cruz, B., Noetel, M., & Lonsdale, C. (2019). Type of screen time moderates effects on outcomes in 4013 children: evidence from the Longitudinal Study of Australian Children. International Journal of Behavioral Nutrition and Physical Activity, 16, 1-10.

Odgers, C. (2018). Smartphones are bad for some teens, not all. Nature, 554(7693), 432-434.

Granic, I., Lobel, A., & Engels, R. C. (2014). The benefits of playing video games. American psychologist, 69(1), 66.

Page, A. S., Cooper, A. R., Griew, P., & Jago, R. (2010). Children's screen viewing is related to psychological difficulties irrespective of physical activity. Pediatrics, 126(5), e1011-e1017.

Babic, M. J., Smith, J. J., Morgan, P. J., Eather, N., Plotnikoff, R. C., & Lubans, D. R. (2017). Longitudinal associations between changes in screen-time and mental health outcomes in adolescents. Mental Health and Physical Activity, 12, 124-131.

Michl, L. C., McLaughlin, K. A., Shepherd, K., & Nolen-Hoeksema, S. (2013). Rumination as a mechanism linking stressful life events to symptoms of depression and anxiety: longitudinal evidence in early adolescents and adults. Journal of abnormal psychology, 122(2), 339.

Millgram, Y., Joormann, J., Huppert, J. D., & Tamir, M. (2015). Sad as a matter of choice? Emotion-regulation goals in depression. Psychological science, 26(8), 1216-1228.

Owens, K. M., Asmundson, G. J., Hadjistavropoulos, T., & Owens, T. J. (2004). Attentional bias toward illness threat in individuals with elevated health anxiety. Cognitive Therapy and Research, 28, 57-66.

Yoshie, M., & Haggard, P. (2013). Negative emotional outcomes attenuate sense of agency over voluntary actions. Current Biology, 23(20), 2028-2032.

Caspi, A., Houts, R. M., Belsky, D. W., Goldman-Mellor, S. J., Harrington, H., Israel, S., ... & Moffitt, T. E. (2014). The p factor: one general psychopathology factor in the structure of psychiatric disorders?. Clinical psychological science, 2(2), 119-137.

Chapter 2

Individual Differences in Information-Seeking

2.1 Overview

Thanks to advances in technology, massive amounts of information are now easily accessible. This includes personalised information about people's past, present and future. Individuals must make many decisions regarding which information they would like to receive and which they would rather avoid. It is unclear how people make these choices.

Despite the relevance of this question to domains such as health, politics and science, we know surprisingly little about what drives information-seeking. Nor do we have a clear understanding of why an individual decides to seek out particular information, while another actively avoids it. For example, a recent study (Sunstein, 2019) found that approximately half of individuals surveyed wanted to know if they had a genetic predisposition to cancer, while the other half did not; half wanted to know the estimated global temperature in 2100, half did not; half wanted to know the amount of calories in meal options, half did not. In this Chapter, we characterise and quantify motives of information-seeking and show how they explain individual-differences in information-seeking choices.

Sharot and Sunstein (2020) have recently proposed a theory which characterises the key motives for information-seeking. According to this theory, when deciding whether to seek information, people first estimate what the information will reveal and then estimate the expected impact of that information on their Action, Affect and Cognition. With regards to Action, the prediction is that people want information more when it can aid in selecting action that will help gain rewards and avoid harm (Sharot & Sunstein, 2020; Stigler, 1961). For example, a person hosting a Thanksgiving dinner would be more likely to want to know how to prepare a turkey for the occasion. With regards to Affect, people will be more likely to want information when they expect knowledge will make them feel better than ignorance (and vice versa; Sharot & Sunstein, 2020; Hertwig & Engel, 2016; Persoskie et al., 2014; Golman et al., 2017; Karlsson et al., 2009; Lerman et al., 1998; Kobayashi et al., 2019; Charpentier et al., 2018; Cuthbert & Insel, 2013). For example, the prediction is that a person would be more likely to want to know the outcome of their favourite sports team game if they believed they had won. With regards to Cognition, people will want information about concepts they think of often (Sharot & Sunstein, 2020). This is because such information is especially relevant to their internal representation of their world and highly connected with many other concepts (Sharot & Sunstein, 2020). For example, the prediction is that those who often think about Ireland would likely want to know the latest developments in Irish news compared to someone who doesn't.

It is also possible that people will seek information to reduce their uncertainty (Afifi & Weiner, 2004; Berlyne, 1957; Kreps & Porteus, 1978) or, conversely, to confirm their beliefs (Nickerson, 1998; Klayman & Ha, 1987; Kappes et al., 2020). Therefore, this factor should also be considered.

The estimated impact of information on action, affect and cognition is referred to as Instrumental Utility, Hedonic Utility and Cognitive Utility, respectively (Sharot & Sunstein, 2020). Each of these estimates can be positive (increasing information seeking), negative (increasing information avoidance) or zero (inducing indifference; Sharot & Sunstein, 2020). We hypothesised that these estimates are integrated into a computation of the value of information, which will trigger information seeking or its active avoidance (Sharot & Sunstein, 2020). Here, over five studies testing 543 participants we provide an empirical test of this theory. To examine if the theory is domain general or domain specific, we test information-seeking in three different domains – information about self-traits, finance and health.

We had further proposed that each of the three factors may be weighted differently, influencing the decision to seek or avoid information to different degrees (Sharot & Sunstein, 2020; **see Figure 1.1a**). Individual differences in information-seeking may be related to the different weight individuals assign to each motive. For example, certain individuals may care most about the Instrumental Utility of information, whereas others may care most about the need to regulate their affective state, while other may assign equal weight to all three motives when seeking information, etc. Here, we quantify those differences and examine to what degree they are stable, or change, over time within and across domains, by conducting three longitudinal studies.

If there are individual differences in the importance people place on different types of information, such differences might be related to well-being (Sharot & Sunstein, 2020). Many psychopathology symptoms can be broadly characterized as problems in affective processes, cognitive functions, and action planning and execution (Cuthbert & Insel, 2010; Cuthbert & Insel, 2013). Thus, abnormalities in these domains may reveal themselves in the type of information people choose to seek or avoid. For instance, poor well-being is characterised by a reduction in the belief that one has agency over outcomes (et al., 2020; Alloy & Abramson, 1979; Yoshie & Haggard, 2013), which may lead to a reduction in the impact of instrumental utility on information-seeking. Additionally, people with poor well-being have been shown to have a negativity bias in the context of information-seeking (Owens et al., 2004; Aderka, et al., 2013), thus in the context of the integrative information-seeking framework, they may be more likely to seek information that they attribute a negative hedonic value to. It’s also possible that the desire to know information frequently thought about can lead to worse mental health or promote resilience due to cognitive closure. Research indicates that being at either extreme of the spectrum of ambiguity tolerance can be detrimental to mental health (Anderson & Schwartz, 1992).

Moreover, the specificity of these motives might reflect a broader susceptibility to mental health challenges. This idea aligns with the *P-factor* hypothesis, which suggests the existence of a single underlying factor contributing to the risk and severity of various mental health disorders (Caspi et al., 2014). As poor mental health is often associated with problems related to self-perception and thoughts regarding the self (Christensen et al., 2003; Graham et al., 2014; Sass & Parnas, 2003; Jacobi et al., 2004; Doron et al., 2008; Silverstone & Salsali, 2003), we test the relationship between mental health and information-seeking in the domain of self-referential knowledge. If indeed psychopathology symptoms are related to specific patterns of information-seeking, there is potential for using measured markers of information-seeking as an assessment to indicate vulnerability to mental health problems.

Given this rich potential, it is surprising how limited our knowledge is of the links between mental health and information-seeking. In fact, despite information-seeking being central to human behaviour, we know remarkably little about how to quantify it or the mechanisms that underlie it. To address these unknowns, we conducted five studies in which participants were asked to indicate whether they would want to receive 40 pieces of information. In Study 1, 2 and 5 the information was related to self-traits, in Study 3 to finance and in Study 4 to health. Participants also provided ratings which served as proxies for the Instrumental, Hedonic and Cognitive Utility they assigned to each potential piece of information. These proxies were then used to quantify participants' information-seeking motives and explain individual differences in participants' choices. Study 1 and 3 were longitudinal studies that enabled us to quantify the stability of the motives over time within an individual and domain, and Study 4 examined stability over time across domains. Additionally, in Studies 1 and 2 we assessed participants' mental health using a battery of self-report psychopathology questionnaires (Foa et al., 2005; Zung, 1965; Spielberger & Gorsuch, 1983; Saunders et al., 1993; Marin et al., 1991; Garner et al., 1982; Patton et al., 1995; Mason et al., 2005; Fresco et al., 2001) and examined these responses for an association between mental health and information seeking motives. In particular, we implemented a dimensionality approach (Gillan et al., 2016; Rouault et al., 2018; Seow & Gillan, 2020), which considers the possibility that a specific symptom is predictive of several psychiatric conditions, thus allowing an investigation that cuts through classic clinical boundaries.

2.2 Methods

Participants (Study 1). Ninety-nine participants completed the task on Mechanical Turk online system. Data of 3 subjects who did not pass the attention checks were excluded from further analysis. In particular, participants were asked five times throughout the Study to select a particular answer (for example: 'Please click answer two'). This is to ensure that participants are being attentive. Participants who answered more than one of the attention checks incorrectly, were excluded from analysis. Of those who passed the check, 16 gave the same exact response on all trials in at least one of the utility ratings and thus their beta coefficients could not be calculated. Thus, data of 80 subjects were analysed (age = 37.69, SD = 9.18; females = 46.3%). One stimulus was repeated twice due to a coding error and thus data of the second repetition was removed from analysis leaving data from 39 trials per subject. Participants received £7.50 for their participation.

Note, for all studies presented in this article, ethical approval has been provided by the Research Ethics Committee at University College London and all participants have given their informed consent to participate.

Procedure (Study 1). Participants were asked to imagine that their family/friends had rated them on different attributes taken from Allport and Odbert (1936). For example, 'intelligent', 'unreliable', (**see Materials** for all attributes). In block one, on each of 40 trials participants indicated whether they would want to know how others rated them on a specific attribute using a six-point Likert scale from (definitely don't want to know) to (definitely want to know) (**Supplementary Figure 2.2**). This was self-paced. Half the attributes were positive and half negative. Traits were presented in a random order.

In block two, to assess Instrumental, Hedonic and Cognitive Utility participants provided the following ratings for each attribute respectively (self-paced): (i) Their expectations regarding how useful each piece of information would be (from -3 'not useful ' to +3 very useful), which provided an estimate of Instrumental Utility (e.g., how useful would it be to know how others rated you on 'intelligence'?); (ii) How they expect to feel if the rating was revealed to them (from -3 'very bad' to +3 'very good') (e.g., how will you feel if you knew how others rated you on 'intelligence'?) and how they expect to feel if the rating was never revealed to them (from -3 'very bad' to +3 'very good'; e.g., how will you feel if you never knew how others rated you on 'intelligence'?). The difference between the last two ratings provided an estimate for Hedonic Utility. This calculation was necessary as the emotional impact of knowing versus not knowing is not simply the inverse of each; for instance, a person might feel equally or more distressed by choosing ignorance over knowledge; (iii) How often they think about each attribute (from -3 'never 'to +3 'very often') (e.g., how often do you think about 'intelligence'?), which provided an estimate of Cognitive Utility. This can be in relation to themselves, to others or to the concept itself. What we are measuring is how often the concept is thought of regardless of the exact context. The questions were selected based on the theory paper (Sharot & Sunstein, 2020) in which we had introduced the three utilities of information-seeking and suggested these quantifiable predictions. We note that these are not necessarily the only questions one can use to measure the three utilities, but we had proposed them as central ones (Sharot & Sunstein, 2020). Participants also indicated how they expected others would rate them (from -3 'not at all this trait' to +3 'very much this trait'; scores were reversed for negative valanced stimuli) and their confidence in this rating (-3 'not certain' to +3 'very certain'). The reason we asked about expectations is that it allowed us to then assess whether people were more likely to seek knowledge when they are confident or unconfident about what the information will reveal. Indeed, many studies suggest that uncertainty is related to information seeking (Afifi & Weiner, 2004; Berlyne, 1957; Kreps & Porteus, 1978; Nickerson, 1998; Klayman & Ha, 1987; Kappes et al., 2020; Stigler, 1961; Kobayashi et al., 2019; Charpentier et al., 2018; Hirshleifer & Riley, 1979; van Lieshout et al., 2018; Trudel et al., 2021). Sometimes people want information about things they are certain about (a form of confirmation bias; Nickerson, 1998; Klayman & Ha, 1987; Kappes et al., 2020) and sometimes they want information about things they are uncertain about (Afifi & Weiner, 2004; Berlyne, 1957; Kreps & Porteus, 1978), with one study suggesting that the sign of the effect can vary according to the environment (Trudel et al., 2021). Each question was displayed separately for each attribute. Descriptive statistics of these ratings and their inter-relationships are displayed in **Supplementary Table 2.1**.

Next, participants completed self-report questionnaires which assess psychopathology symptoms (Foa et al., 2005; Zung, 1965; Spielberger & Gorsuch, 1983; Saunders et al., 1993; Marin et al., 1991; Garner et al., 1982; Patton et al., 1995; Mason et al., 2005; Fresco et al., 2001), the list is adapted from Gillan and colleagues (2016). These included: Obsessive-Compulsive Inventory – Revised (OCI-R; Foa et al., 2005), Self-Rating Depression Scale (SDS; Zung, 1965), State-Trait Anxiety Inventory (STAI; Spielberger & Gorsuch, 1983), Alcohol Use Disorder Identification Test (AUDIT; Saunders et al., 1993), Apathy Evaluation Scale (AES; Marin et al., 1991), Eating Attitudes Test (EAT-26; Garner et al., 1982), Barratt Impulsivity Scale (BIS-11; Patton et al., 1995), Short Scales for Measuring Schizotypy (Mason et al., 2005), Liebowitz Social Anxiety Scale (LSAS; Fresco et al., 2001). Participants also indicate their age, gender, annual income and level of education. The task was coded using the Qualtrics online platform (https://www.qualtrics.com). Analysis was conducted using IBM SPSS 27 and R studio (Version 1.3.1056). All statistical tests conducted in the present article are two-sided.

Materials (Study 1). The following traits were used (adapted from Allport & Odbert, 1936): Courageous, Shy, Honest, Enthusiastic, Lazy, Mean, Trustworthy, Cooperative, Self-centered, Generous, Incompetent, Considerate, Rude, Conscientious, Boring, Easy-Going, Carless, Curious, Sophisticated, Unhelpful, Cowardly, Deceitful, Sociable, Confident, Unmotivated, Unfriendly, Unreliable, Organized, Greedy, Selfish, Polite, Disorganized, Imaginative, Adaptable, Ignorant, Competent, Immature, Helpful, Narrow-minded, Kind.

Model Testing (Study 1). We first tested the prediction that information-seeking choices across participants are best explained considering Instrumental Utility, Hedonic Utility and Cognitive Utility. To that end we ran a general linear mixed-effects model to assess the effect of the three utilities on information-seeking choice. The dependent variable was choice which is defined as the rating a participant gave to the question how much they wanted to know how others rated them on the respective trait. We quantified the scale such that one end ('definitely don't want to know') was given a -3 and the other ('definitely want to know') a +3. The three predictors were (i) Instrumental Utility (i.e., participants' rating of how useful it would be to receive that piece of information), (ii) Hedonic Utility (i.e., participants' rating of how they would feel if they received information minus how they would feel if they remain ignorant), (iii) Cognitive Utility (i.e., participants' rating of how often they think of the concept).

Each of these three factors was mean centred within-participant and rating across all trials before entering in the model as fixed effects and random effects. Random intercepts and slopes were included for each participant as well as random intercepts for each item. This model (model 1) is the 'hypothesised model'. Six comparison models were tested which included only one or two utilities each. We compared the Bayesian Information Criterion (BIC; Schwarz, 1978) and the Akaike information criterion (AIC; Bozdogan, 1987) scores of all seven models (the full hypothesised model and six comparison models) to test whether the full hypothesised model fits best. The BIC and AIC penalises models for complexity (Schwarz, 1978; Bozdogan, 1987. We also attempted to include a random slope for each item (Barr et al., 2013; Judd et al., 2012; Murayama et al., 2014), however the theorised and comparison models frequently failed to converge across studies. Thus, in line with recommendations (Matuschek et al., 2017; Eager & Roy, 2017), we reduced the item random effect structure (taking away the item random slope) which successfully improved the convergence of models. Importantly, we did not observe any difference in the significance of the predictors between the model structures for the times that the model was able to converge.

Additional comparison models examined whether adding participants' confidence regarding what the information will reveal provided a better fit. In particular, we added a fourth factor to the full model: participants' rating of how confident they are of how others will rate them (again mean centred within participant). We compared the BIC and AIC scores of that model to the original hypothesised model, which only includes the three factors. We also tested models including subsets of those four factors that include the confidence rating (i.e., all models including only three factor or two factors, where one of the factors is the confidence rating and a model that includes only confidence ratings) to see whether any provide a better fit to the data than our hypothesised three factor model. The winning model (i.e., model with lowest BIC and AIC score) was used for all the analyses below.

Relating Information-Seeking Types to Mental Health (Study 1). Each participant was scored on the three psychopathology dimensions identified by Gillan and colleagues (2016) and replicated by Rouault and colleagues (2018) 'Anxious-Depression', 'Social-Withdrawal' and 'Compulsive-Behaviour and Intrusive Thought'. To generate these scores, we first Z-scored the ratings for each questionnaire item separately across subjects. Next, we multiplied each Z-scored item by its factor weight as identified earlier (**Figure 2.2a**). Then for each subject the three psychopathology dimension scores were calculated by summing all of the weighted items assigned to each dimension. Nine participants did not compete all questionnaires and therefore were not included in the mental health analysis.

For each participant a general linear model was conducted predicting information choice on each trial from the three utilities. This generated three beta coefficients, indicating the weight each participant assigned to each motive when seeking information. These were then related to the psychopathology dimensions by submitting the three psychopathology dimension scores into a mixed ANOVA with psychopathology dimension as a within-subject factor and Instrumental Utility (β1), Hedonic Utility (β2), and Cognitive Utility (β3) each as within subject modulating covariates as well as participants' age and gender as between subjects modulating covariates. This analysis was then followed up with a simplified analysis in which the average of the three psychopathology scores of each individual were entered as a dependent measure in a linear regression with each of the three beta coefficients (the weight put on Instrumental Utility (β1), Hedonic Utility (β2), and Cognitive Utility (β3)) entered as an independent measure as well as age and gender.

We report whether the three betas reflecting the weight each subject assigned to each information-seeking motive (β1, β2, β3) relate to demographics (age, gender, and education), information-seeking choice, utility ratings, expected information and confidence in this estimation and scores on individual psychopathology questionnaires (Foa et al., 2005; Zung, 1965; Spielberger & Gorsuch, 1983; Saunders et al., 1993; Marin et al., 1991; Garner et al., 1982; Patton et al., 1995; Mason et al., 2005; Fresco et al., 2001) by submitting each into a one-way ANOVA. All significant results were followed up with post-hoc pairwise comparisons. Psychopathology questionnaire scores were corrected for multiple comparisons across nine questionnaires using Bonferroni correction.

We also correlated each of the three-psychopathology dimension scores separately with: information-seeking choice, utility ratings, expected information and confidence in this estimation.

Participants (Study 2). 200 participants completed at Time 1 the same exact task as in Study 1 on Prolific's online platform. All participants who passed the attention check and for whom we could calculate all beta coefficients (i.e., those who did not give the same exact response on all trials in at least one of the utility ratings) (N= 176; age = 28.00, SD = 9.66; females = 47.2%) were then invited to complete the task again three weeks later (Time 2). Out of those, 137 participants completed the task at Time 2, of which 124 participants passed the attention check and did not give the same exact response on all trials in at least one of the utility ratings (age = 26.93, SD = 8.30; females = 46.0%). At Time 1, one random attribute was not presented to each participant due to a coding error, leaving 39 of the 40 attributes to be analysed. At Time 2, participants saw 40 new attributes. Participants received £7.50 for their participation at Time 1 and £3.25 at Time 2.

Procedure (Study 2). At Time 1, participants were asked to complete the exact same procedure as Study 1, outlined previously. Three weeks later (Time 2), participants were asked to complete the same information-seeking task but with 40 different attributes (Allport and Odbert, 1936; see below). Descriptive statistics of these ratings and their inter-relationships are displayed in **Supplementary Tables 2.2 & 2.3**.

Materials (Study 2). In Time 1 we used the same traits (Allport and Odbert, 1936) as in Study 1. In Time 2 we used the following traits (Allport and Odbert, 1936): Open-minded, Intelligent, Objective, Admirable, Calm, Loyal, Humble, Disciplined, Efficient, Fair, Stable, Warm, Wise, Impressive, Gracious, Patient, Popular, Creative, Ambitious, Dedicated, Cruel, Indecisive, Naïve, Disruptive, Reserved, Aggressive, Foolish, Cold, Difficult, Disloyal, Shallow, Messy, Thoughtless, Insensitive, Weak, Impulsive, Fearful, False, Dull, Arrogant.

Analysis (Study 2). We analysed the data from Time 1 exactly as in Study 1. This allowed us to examine for replication of the results of Study 1 and provided us with the three beta coefficients (relating the three motives to information-seeking) for each participant in Time 1.

Next, we examined whether the relative importance of the three information-seeking motives are stable over time within individuals. To do this, we first calculated for each participant the three beta coefficients (relating the three motives for information-seeking) from Time 1 and Time 2 data separately. Then we measured by how much each participant moved over time with respect to each of their 3 motives, with each beta coefficient indicated on a separate axis in a three-dimensional space. AB denotes the distance between participants at Time 1 and Time 2 in a 3-dimensional space, with each axis representing the weight they place on Instrumental Utility (x-axis), Hedonic Utility (y-axis), and Cognitive Utility (z-axis).

$$AB = \sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2 + (z_2 - z_1)^2} \quad (1)$$

x_2 denotes participants' Instrumental Utility beta at Time 2, x_1 denotes its beta at Time 1, y_2 denotes participants' Hedonic Utility beta at Time 2, y_1 denotes the beta for Hedonic Utility at Time 1, z_2 denotes participants' Cognitive Utility beta at Time 2, while z_1 denotes the beta for Cognitive Utility at Time 1. If the relative weight individuals place on the motives for information-seeking are stable overtime, we would expect this change to be significantly less than what would be expected by chance. To test this, we reran the exact same analysis above for each subject, but each time randomly mismatching one participant's T1 data with another participant's T2 data (i.e., permutation test). We then compared the average distance participants actually moved in the three-dimensional space from T1 to T2 to the average distance calculated from the permutation test. We did this 10,000 times and calculated the percentage of the times the average distance participants actually moved from T1 to T2 was smaller than chance.

We also calculated an Intraclass Correlation Coefficient for each relative weight individuals placed on each of the three motives when seeking information across Time 1 and Time 2. To do this, we mean centred the three betas for each participant and time and then conducted a separate ICC test for each pair of equivalent betas.

When examining whether the motives for information-seeking were related to mental health in Study 2, we implemented the same procedure as in Study 1, entering the average betas across the two time points into all analyses.

Participants (Study 3). One hundred forty-nine participants completed the Study at Time 1 on Prolific's online platform. All participants who passed the attention check and for whom we could calculate all beta coefficients (i.e., those who did not give the same exact response on all trials in at least one of the utility ratings) (N = 122; mean age = 31.91, SD = 9.76 females = 46.7%) were invited to complete the task again three weeks later (Time 2). Out of those, 95 participants completed the task at Time 2. Two participants were not included due to providing different Prolific IDs for each time point. Eighty-two participants (mean age = 32.88, SD = 9.86; females = 52.4%) passed the attention check and did not give the same exact response on all trials in at least one of the utility ratings. Participants that passed the attention checks received £3.25 for their participation at Time 1 and for Time 2.

Procedure (Study 3). Participants were asked to imagine that we possessed a crystal ball that could reveal the answer to any question. In block one, on each of 40 trials they were asked whether they wanted to know specific information related to finance (e.g., what the exchange rate was between Dollar and Pound, what income percentile they fall into etc., see Supplementary Information for all stimuli). On each trial the stimuli were different and differed between Time 1 and Time 2. They indicated their response using a six-point Likert scale from -3 (definitely don't want to know) to +3 (definitely want to know) (**Supplementary Figure 2.2**). This was self-paced.

In block two, participants provided the following ratings for each of the 40 traits: (i) their expectations regarding how useful it would be to know the information (from -3 'not useful' to +3 'very useful'), which provided an estimate of Instrumental Utility (e.g., 'how useful would it be to know X?'); (ii) How they expect to feel if they knew the information (from -3 'very bad' to +3 'very good'; e.g., 'how will you feel if you knew X?') and how they expect to feel if they never knew the information (from -3 'very bad' to +3 'very good'; e.g., 'how will you feel if you never knew how X?'). The difference between the last two ratings provided an estimate for Hedonic Utility; and (iii) how often they think about each topic (from -3 'never' to +3 'very often'; e.g., 'how often do you think about X?'), which provided an estimate of Cognitive Utility. Participants also indicated what they expected the information would be ('what do you think the answer is?'). Depending on the question asked, participants either answered on a scale (e.g., for the question about what the Gross Domestic Profit is, the scale went from 'low' to 'high') or input their answer into a text box (e.g., for the question about what your daily expenses are). Finally, participants indicated their confidence in what they expected the information would reveal (from -3 'not certain' to +3 'very certain').

We highlight a qualitative difference between the expectations scale in Study 1 and 2 and that in Study 3. In studies 1 and 2 participants indicated how they expected others would rate them (from 'not at all this trait' to 'very much this trait'). For the analysis scores were reversed for negative valanced stimuli (e.g., boring). Once expectations for negative valanced stimuli are reversed, this measure tell us how good or bad a subject expects information to be. For example, for 'intelligence' a high rating will indicate a subject believed other saw him/her as possessing this trait (which is a good thing) for 'boring' a low rating will indicate a subject believed other saw him/her as not possessing this trait (which is a good thing). Thus, one could use expectations in these studies in a model where the motive for information is learning good news. However, in Study 3, expectations regarding financial information do not clearly reflect expectations of valence or feelings. If a subject expects the Dollar to Pound exchange rate to be high that does not tell us how they expect to feel if they learn it is high. In fact, there is no clear way to quantify expectations in the financial task nor would there be a consistent way to do so in other tasks like general knowledge questions (e.g., 'Do you want to know if dogs are related to wolfs'). To build a model of motives of information-seeking that can generalise to other domains any of the three utilities (+ confidence) would be possible to include, but not one that includes expectations of the information to be revealed. Descriptive statistics of all ratings and their inter-relationships are displayed in **Supplementary Tables 2.4 & 2.5**.

Analysis (Study 3). We carried out the exact analysis as described in Study 2 to examine whether the three motives are significant predictors of information-seeking and whether the three-factor model is a better fit to the data than other models. We also describe individual differences in the same way and examine stability of weighting of information seeking motives over time as done in Study 2.

Participants (Study 4). We invited all participants who completed Study 3, Time 1 and an additional 101 new participants to take part in this study, which was run on Prolific's online platform. Data of the 116 participants who completed the study, passed the attention check and for whom we could calculate all beta coefficients (i.e., those who did not give the same exact response on all trials in at least one of the utility ratings) was analysed (mean age = 31.15, SD = 11.30, females = 56.9%). Thirty-eight of these are participants who also completed Study 3, Time 1. Participants that passed the attention checks received £5.00 for their participation.

Procedure (Study 4). Participants were asked to imagine that we had information about their genetic makeup. In block one, on each of 40 trials they were asked whether they wanted to know whether or not they carried a gene that increases their likelihood of a particular health condition or trait (e.g., 'Would you like to know if you have a gene that increases your likelihood of Alzheimer's disease?', 'Would you like to know if you have a gene that increases your likelihood of a Strong Immune System?'; **see Supplementary Chapter 2** for all stimuli). On each trial the stimulus was different. They indicated their response using a six-point Likert scale from -3 ('definitely don't want to know') to +3 ('definitely want to know'; **see Supplementary Figure 2.2**). This was self-paced.

In block two, participants provided the following ratings for each of the 40 health condition or traits: (i) their expectations regarding how useful it would be to know the information (from -3 'not useful' to +3 'very useful'), which provided an estimate of Instrumental Utility (e.g., 'how useful would it be to know X?'); (ii) How they expect to feel if they knew the information (from -3 'very bad' to +3 'very good'; e.g., 'how will you feel if you knew X?') and how they expect to feel if they never knew the information (from -3 'very bad' to +3 'very good'; e.g., 'how will you feel if you never knew X?'). The difference between the last two ratings provided an estimate for Hedonic Utility; and (iii) how often they think about each topic (from -3 'never' to +3 'very often'; e.g., 'how often do you think about X?'), which provided an estimate of Cognitive Utility. Participants also indicated their expectations of how likely it is that they carry the gene (from -3 'not likely' to +3 'very likely', e.g., 'how likely is it that you carry this gene?'; scores were reversed for negative valanced stimuli). Finally, participants indicated their confidence in what they expected the information would reveal (from -3 'not certain' to +3 'very certain'). Descriptive statistics of all ratings and their inter-relationships are displayed in **Supplementary Table 2.6**.

Analysis (Study 4). We carried out the exact analysis as described in Study 1, 2 and 3, to examine whether the three motives are significant predictors of information-seeking in the health domain and whether the three-factor model is a better fit to the data than other models. We also describe individual differences in the same way as Study 2 and 3, however, here we examine the stability of the weights given to the motives of information-seeking across time and domains (i.e., finance Study 3, Time1 and health Study 4).

2.3 Results

Information-Seeking is best explained by taking into account Instrumental, Hedonic and Cognitive Utilities (Study 1). We tested 99 participants on the information-seeking task described above. Eighty participants passed the attention check and had enough variability in their rating data to generate three beta coefficients (that is did not insert the same rating for all stimuli on any of the scales). We submitted their data into a mixed-effects model to estimate the relationship between Instrumental Utility, Hedonic Utility and Cognitive Utility (which were estimated using the ratings as described above) and the desire to receive information (**see Methods**). Each of these three factors were centred within-participant for each rating across all trials and included in the model as fixed and random effects. Random intercept and slope were estimated for each participant as well as random intercept for each item (see methods). This revealed a significant fixed effect of Instrumental Utility ($\beta = 0.114 \pm 0.029$ (SE), $t(60.17) = 3.918$, $p < 0.001$, **Figure 2.1b**), Hedonic Utility ($\beta = 0.123 \pm 0.022$ (SE), $t(61.28) = 5.531$, $p < 0.001$, **Figure 2.1b**) and Cognitive Utility ($\beta = 0.091 \pm 0.031$ (SE), $t(89.98) = 2.935$, $p = 0.004$, **Figure 2.1b**). In particular, participants expressed a greater desire for knowledge when they believed the information would be useful, would have a more positive impact on their affect than ignorance, and also for stimuli they thought of frequently (**see Supplementary Chapter 2** for a study testing three additional motives of information seeking). On average participants rated their desire to receive information as 0.43 (SD = 1.30), which is significantly different from the mid-point of the scale, $t(79)=2.970$, $p = 0.004$.

We tested thirteen additional models to test if any account for information seeking choices better than the hypothesised model. These included models in which only a subset of the three utilities were entered and also models including how confident participants were regarding the information to be revealed, which they also provided as a rating. The hypothesised model, which included instrumental, hedonic and cognitive utilities as predictors of information-seeking, fit the data better than all other thirteen models. This is indicated both by a lower Bayesian Information Criterion (BIC) (**Figure 2.1c**) and Akaike information criterion (AIC) (**Supplementary Table 2.8**), both of which penalises models for complexity.

While across our sample all three motives (action, affect, cognition) were strongly associated with information seeking, there may be significant individual differences in the importance participants assign to these when seeking information. To characterise such differences, we conducted for each participant separately a general linear model predicting information choice on each trial from the three utilities. As can be observe in **Figure 2.1d** there were large individual differences in the weight participants assign to each motive. Most of the participants had a dominant motive; over one third of participants (34.75%) assigned more than twice the weight to one utility relative to the other two, and most participants (73.75%) assigned at least 1.25 times more weight to one utility than the other two. Different motives were dominant for different individuals, with action being dominant for 20% of individuals in this sample, affect for 27.5%, cognition for 26.25% and 26.25% did not have one particularly strong motive (that is no motive was assigned a weight at least 1.25 greater than the rest).

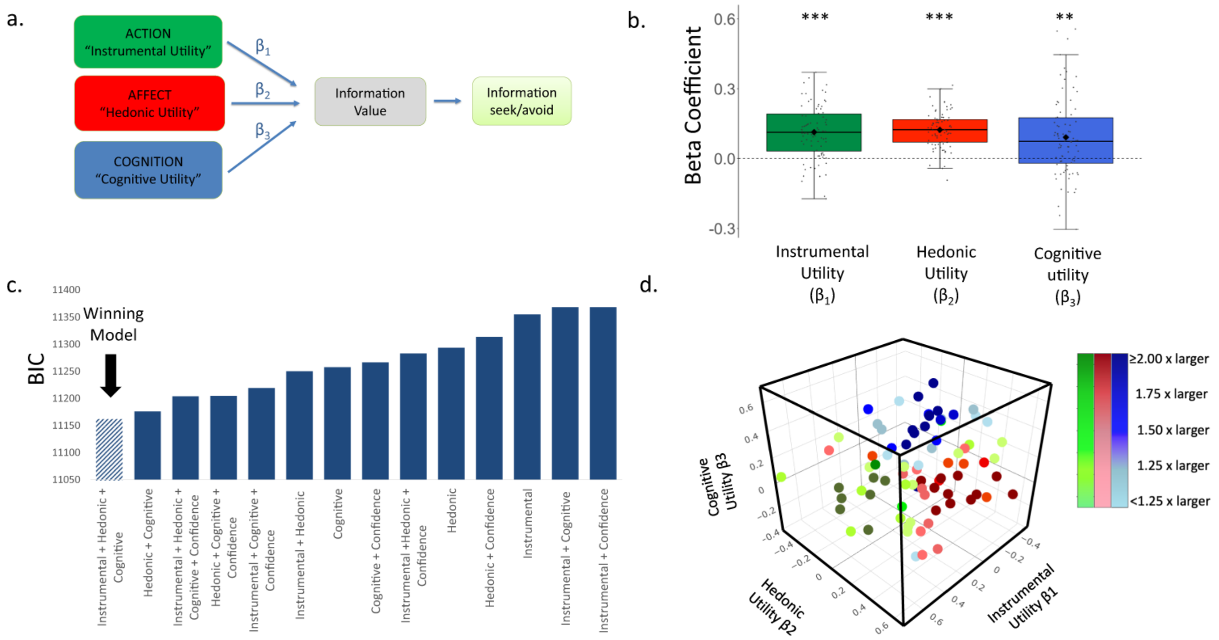

Figure 2.1. Information-Seeking Motives. (a) Information seeking and its avoidance is hypothesised to be driven by Instrumental Utility, Hedonic Utility and Cognitive Utility (Sharot & Sunstein, 2020). These values reflect the predicted impact of information on action, affect and cognition, respectively. These estimates are hypothesised to be integrated into a computation of the value of information, with different weights (β1–3) assigned to each of the three factors. The integrated value can lead to information seeking or avoidance. **(b)** Plotted are the beta coefficients from a linear mixed-effects model (N = 80 participants), showing that participants' desire to receive information was greater when the Instrumental Utility ($p<0.001$, two sided), Hedonic Utility ($p<0.001$, two sided) and Cognitive Utility ($p=0.004$, two sided) of information were higher. These were estimated respectively by participants' ratings of how useful the information would be, how they would feel to know vs not to know, and how frequently they think about the stimulus. The horizontal lines indicate median values, boxes indicate 25–75% interquartile range and whiskers indicate 1.5 × interquartile range; individual scores are shown as dots. **(c)** BIC scores reveal that the model described in b fit the data better than models including alternate combinations of the utilities and also those including participants' confidence regarding what the information would reveal. The same was true when examining AIC scores (**see Supplementary Table 2.8**). Smaller BIC and AIC scores indicate better fit. **(d)** Plotted are the weights each individual put on each motive when seeking information. Beta coefficients of Instrumental Utility are on the x-axis, of Cognitive Utility on the y-axis and of Hedonic Utility on the z-axis. Green dots represent participants who put the largest weight on Instrumental Utility when seeking information. Red dots represent participants who put the largest weight on Hedonic Utility when seeking information. Blue dots represent

participants who put the largest weight on Cognitive Utility when seeking information. The colour gradient represents how dominant the largest weight was in comparison to the other two weights. Individuals who put more than twice as much weight on their dominant utility than the other two utilities are represented in darkest colours. Those whose dominant utility was less than 1.25 times larger than the other two are represented in the lightest colours. ***$P < 0.001$, **$P < 0.01$ (two sided).

Individual Differences in the Weights Assigned to Information-Seeking Motives Provide a Window into Mental Health (Study 1). As described in the introduction, our hypothesis was that the different weights individual assigned to the different motives were related to mental health. Thus, we tested for a relationship between beta coefficients across individuals and mental health. We measured mental health using a dimensionality approach (Gillan et al., 2016; Rouault et al., 2018; Seow & Gillan, 2020). This approach considers the possibility that a specific psychopathology symptom is predictive of several conditions, allowing an investigation that cuts through classic clinical psychopathology boundaries. In particular, previous work (Gillan et al., 2016; Rouault et al., 2018; Seow & Gillan, 2020) used a factor analysis across items in a large battery of traditional psychopathology questionnaires (Foa et al., 2005; Zung, 1965; Spielberger & Gorsuch, 1983; Saunders et al., 1993; Marin et al., 1991; Garner et al., 1982; Patton et al., 1995; Mason et al., 2005; Fresco et al., 2001) and identified three psychopathology dimensions (Gillan et al., 2016; Rouault et al., 2018) across those items: 'Anxious-Depression', 'Social-Withdrawal' and 'Compulsive-Behaviour and Intrusive Thought'. The factor analysis provided a weight to each item in relation to each dimension (**Figure 2.2a**). Thus, a person's symptom severity for each dimension can be quantified by having an individual complete a battery of traditional psychopathology questionnaires (Foa et al., 2005; Zung, 1965; Spielberger & Gorsuch, 1983; Saunders et al., 1993; Marin et al., 1991; Garner et al., 1982; Patton et al., 1995; Mason et al., 2005; Fresco et al., 2001) and then calculating a weighted average across items' ratings. Indeed, this is what we did for each participant. First, we Z-scored the ratings of each questionnaire item separately across participants (not Z-scoring does not alter the significance of results). Then, for each participant we calculated the three-dimension scores which we submitted into a mixed ANOVA with psychopathology scores ('Anxious-Depression', 'Social-Withdrawal', 'Compulsive-Behaviour and Intrusive Thought') indicated as a within-subjects factor and the weight put on Instrumental Utility ($\beta1$), Hedonic Utility ($\beta2$), and Cognitive Utility ($\beta3$) when seeking information all indicating within-subject modulating covariates. Participants' age and gender indicated between-subject modulating covariates. We observed a significant main effect of Cognitive Utility on psychopathology scores ($F(1,65) = 6.061$, $p = 0.016$, partial eta square = 0.085).There were no significant effects of Instrumental Utility ($F(1,65) = 2.882$, $p = 0.094$, partial eta square = 0.042) or Hedonic Utility ($F(1,65) = 0.027$, $p = 0.870$, partial eta square = 0.000). No other effects or interactions were significant (all P's > 0.188). These results suggest that the weight participants' assign to Cognitive Utility, but not the other two utilities, when seeking information is related to their mental health across the three psychopathology dimensions, with greater weight on Cognitive Utility associate with better mental health.

To illustrate this result in a more simplified manner, we conducted a linear regression with mental health as the dependent measure (quantified as the average psychopathology score across the three dimensions) and the following predictors: the weight assigned to Instrumental Utility (β1) when seeking information, as well as that assigned to Hedonic Utility (β2) and to Cognitive Utility (β3). Age and gender were also included as predictors. Confirming the analysis above, a significant inverse relationship was observed between mental health and the weight assigned to Cognitive Utility when seeking information (β = -1.053, p = 0.016), suggesting that participants who seek information more on issues they think of often are the ones who report less psychopathology symptoms across the board. No other predictor was significant (Instrumental Utility: β = -0.710, p = 0.094; Hedonic Utility: β = -0.072, p = 0.870; Age: β = -0.010, p = 0.893; Gender: β = -0.211, p = 0.296; **Figure 2.2b**). Finally, correlating each beta with the average psychopathology score across participants (controlling for age and gender), again reveals a significant association with the weight assigned to Cognitive Utility when seeking information (r = -0.244 (67) p = 0.043), but not with the weight assigned to Instrumental (r = -0.136 (67) p = 0.264) or Hedonic (r = 0.09 (67), p =0.463) utilities.

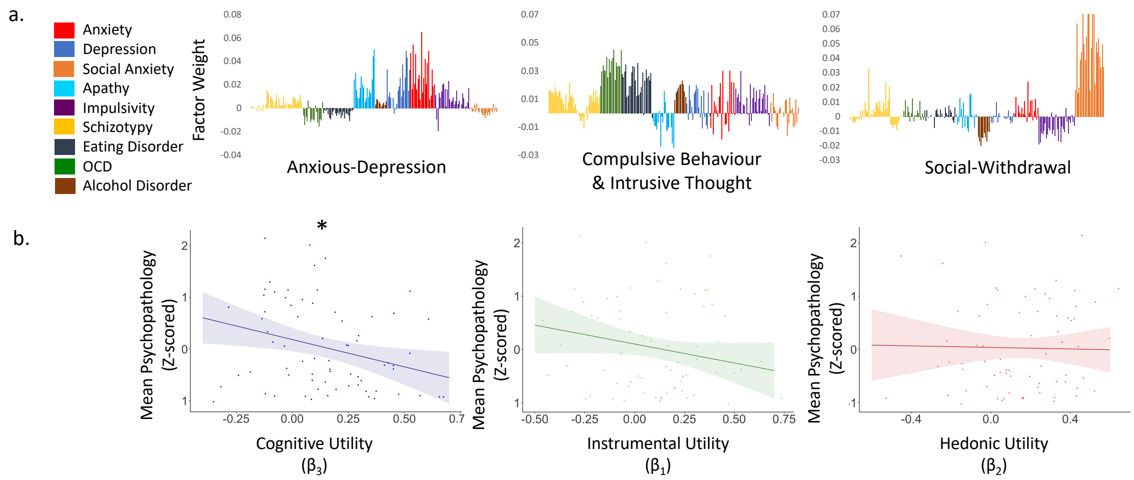

Figure 2.2. Information-Seeking related to Psychopathology. (a) Plotted are the weights, based on Gillan and colleagues (2016), given to each questionnaire item when calculating the weighted score for each subject on each of the three psychopathology dimensions identified previously (‘Anxious-Depression’, ‘Social-Withdrawal’ and ‘Compulsive-Behaviour and Intrusive Thought’). **(b)** Plotted on the y-axis is the average psychopathology score across the three dimensions described in **(a)**, Z-scored. On the x-axis are the weights assigned to each information-seeking motive from a linear regression predicting information-seeking from Instrumental Utility (green), Hedonic Utility (red) and Cognitive Utility (blue). Dots represent individual subjects. Shading represents confidence interval. Line represents the relationship between the abscissa and ordinate controlling for the effect of the other two motives as well as of age and gender. As can be observed, participants who placed a large positive weight on Cognitive Utility when seeking information reported less

psychopathology symptoms ($p = 0.016$, two-sided), while we observed no effect of Instrumental Utility ($p = 0.094$, two-sided) or Hedonic Utility ($p = 0.870$, two-sided). Error bars SEM. *$P < 0.05$ (two-sided). N = 71 subjects.

Stability of information-seeking motives over time (Study 2). Thus far, we have shown that the weights individuals place on motives for information-seeking are meaningful as they provide a window into mental health, which is known to be a function of both of 'trait' and 'state'. If information-seeking styles reflect mental health, they too may be a function of 'trait' and 'state'. One may thus predict that the weights assigned to information-seeking motives may show some stability over time, which also allows for changes due to factors such as altering mood, environment etc.

To quantify the stability of the motive weights of information-seeking over time, we conducted a second longitudinal, study. This study also provided a replication test for the results obtained in Study 1. We tested 200 participants on the same information-seeking task as described above (Time 1), of which 176 participants passed attention checks and had enough variability in their rating data (that is did not insert the same rating for all stimuli on any of the scales) to generate three beta coefficients. Three weeks later we contacted these participants again, inviting them to participate in a follow up study (Time 2). 137 completed the follow up study, on average 22 days following Time 1. Of these, 124 participants passed attention checks and had enough variability in their rating data to generate three beta coefficients. The task at Time 2 was identical to Time 1 except that we used a different list of attributes. This design allowed us to test how stable the relative importance of the three motives of information seeking were over time and stimuli sets. Descriptive statistics of ratings and their inter-relationships are displayed for Time 1 in **Supplementary Table 2.2** and Time 2 in **Supplementary Table 2.3**.

Analysis was conducted as in Study 1. We observed a significant fixed effect of Instrumental Utility (Time 1: $\beta = 0.078 \pm 0.018$ (SE), $t(160.53) = 4.382$, $p < 0.001$, **Figure 2.3a**; Time 2: ($\beta = 0.086 \pm 0.020$ (SE), $t(87.56) = 4.267$, $p < 0.001$, **Figure 2.3c**), Hedonic Utility (Time 1: $\beta = 0.104 \pm 0.016$ (SE), $t(139.18) = 6.348$, $p < 0.001$, **Figure 2.3a**; Time 2: $\beta = 0.135 \pm 0.019$ (SE), $t(90.66) = 7.245$, $p < 0.001$, **Figure 2.3c**) and Cognitive Utility (Time 1: $\beta = 0.050 \pm 0.015$ (SE), $t(173.50) = 3.298$, $p < 0.001$, **Figure 2.3a**; Time 2: $\beta = 0.085 \pm 0.019$ (SE), $t(124.76) = 4.500$, $p < 0.001$, **Figure 2.3c**). As in Study 1, at Time 1 and Time 2 the model which included Instrumental, Hedonic and Cognitive Utilities as predictors of information-seeking, fit the data better than comparison models according to the AIC score (**see Supplementary Table 2.8**). This was also true at Time 1 according to the BIC score (**see Figure 2.3b**), while at Time 2 this model was second best, with a simpler model without Instrumental Utility receiving a lower BIC score. We suggest caution in interpreting this specific score as evidence against the importance for instrumental utility, as this conclusion will go against the AIC result, which penalises less for complexity, as well as all other BIC results in all studies described in this study (Study 1, Study 2 Time 1, Study 3 Time 1, Study 3 Time 2, Study 4). Note, that at Time 1 one competing model (confidence + the three factors) did not converge.

Once again, most individuals had a dominant motive. 43.75% of individuals assigned more than twice the weight to one motive than the other two at Time 1 and 44.35% at Time 2, and 81.18% of individuals assigned at least 1.25 times more weight to one motive than the other two at Time 1 and 81.45% at Time 2. Different motives were dominant for different individuals, with action being dominant for 25.57% of individuals at Time 1 and 20.97% at Time 2, affect for 32.95% at Time 1 and 30.65% at Time 2, cognition for 23.30% at Time 1 and 29.84% at Time 2, and 18.18% at Time and 18.55% at Time 2 did not have one particularly strong motive (that is no motive was assigned a weight at least 1.25 greater than the rest).

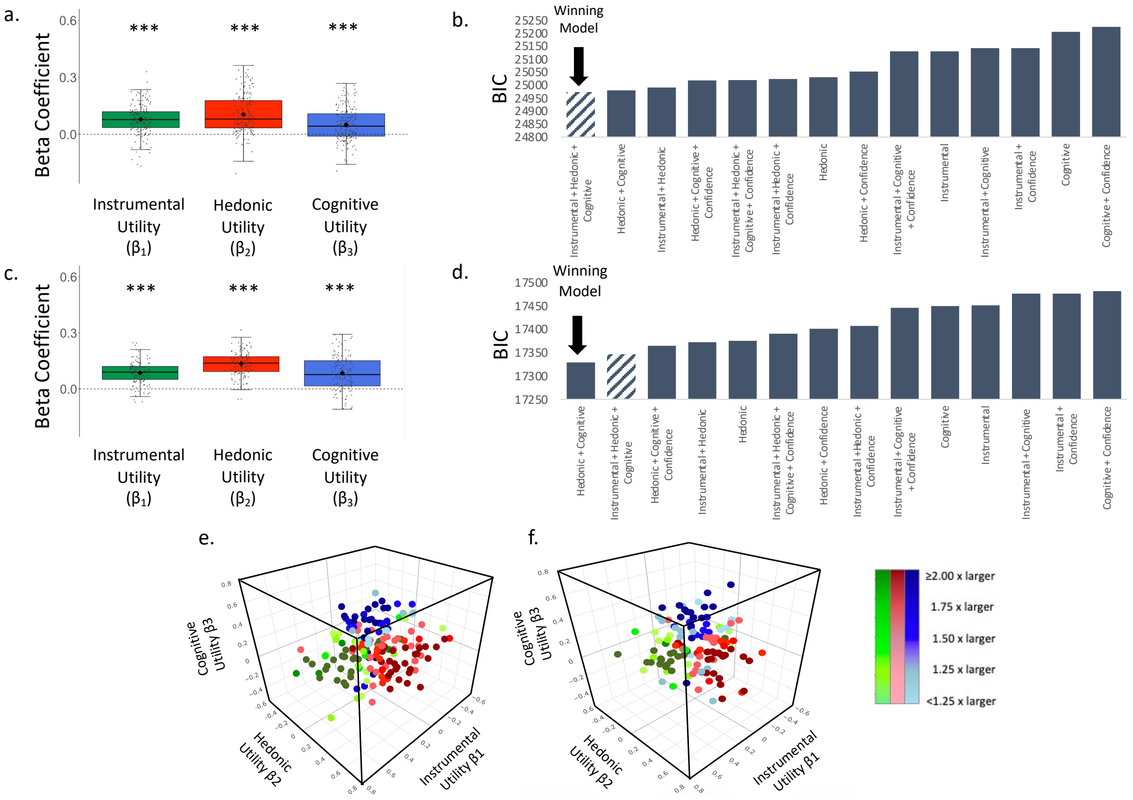

Figure 2.3. Information-Seeking Motives, Study 2. (a&c) Plotted is a boxplot depicting the beta coefficients from a linear mixed effects model at **(a)** Time 1 (N = 176 subjects) and **(c)** Time 2 (N = 124 subjects), which shows that subjects' desire to receive information was greater when the Instrumental Utility (Time 1 $p < 0.001$, Time 2 $p < 0.001$; two-sided), Hedonic Utility (Time 1 $p < 0.001$, Time 2 $p < 0.001$; two-sided) and Cognitive Utility (Time 1 $p < 0.001$, Time 2 $p < 0.001$; two-sided) of information were higher. These were estimated respectively by subjects' ratings of how useful the information would be, how they would feel to know vs not know, and how frequently they think about the stimulus. For each boxplot, the horizontal lines indicate median values, boxes indicate 25–75% interquartile range and whiskers indicate 1.5 × interquartile range; individual scores are shown separately as dots. **(b)** BIC scores from Time 1 reveal that the model described in **(a)** fit the data better than models including other combinations of the utilities and those including subjects' confidence regarding what the information would reveal. **(d)** For Time 2 the model described in **(c)** fit the data second best according to the BIC model. AIC values (reported in **Supplementary Table 2.8**), however, indicate that the model described in **(a&d)** did fit the data best in comparison to control models for Time 1 and Time 2.

Smaller BIC and AIC scores indicate better fit. **(e&f)** Plotted are the weights each individual put on each motive when seeking information at **(e)** Time 1 and **(f)** Time 2. Beta coefficients of Instrumental Utility are on the x-axis, of Cognitive Utility on the y-axis and of Hedonic Utility on the z-axis. Green dots represent participants who put the largest weight on Instrumental Utility when seeking information. Red dots represent participants who put the largest weight on Hedonic Utility when seeking information. Blue dots represent participants who put the largest weight on Cognitive Utility when seeking information. The colour gradient represents how dominant the largest weight was in comparison to the other two weights. Individuals who put more than twice as much weight on their dominant utility than the other two utilities are represented in darkest colours. Those whose dominant utility was less than 1.25 times larger than the other two are represented in the lightest colours. *** = $P < 0.001$ (two-sided).

We next tested to what extent the relative importance of the three information-seeking motives are stable over time within individuals. First, we measured by how much each participant moved over time within the three-dimensional space plotted in **Figures 2.3e&f**. This indicates changes in the relative weights a participant assigned to the three betas. We then tested whether the magnitude of that change was significantly smaller than chance. To test this, we reran the exact same analysis above for each subject, but each time mismatching one participant's T1 data with another participant's T2 data (i.e., permutation test). We then compared the average distance participants actually moved in the three-dimensional space from T1 to T2 to the average distance calculated from the permutation test. We did this 10,000 times and found that 100% of the time the average distance participants actually moved from T1 to T2 was smaller than chance (mean difference between iterations and actual mean movement = 0.103, range of differences = 0.04-0.157). Second, we calculated the Intraclass Correlation Coefficient (ICC) of each beta type across time (**see Methods**). The ICC for each of beta type across time was significant (Instrumental Utility: ICC = 0.302, $p < 0.001$; Hedonic Utility: 0.543, $p < 0.001$; and Cognitive Utility: 0.560, $p < 0.001$).

The relationship between information-seeking and mental health is robust to replication (Study 2). We next examined whether the three motives for information-seeking were related to mental health in Study 2. To do so we calculated each participants' scores on the three psychopathology dimensions (Gillan et al., 2016; Rouault et al., 2018) as indicated in Study 1 and entered these into a mixed ANOVA with psychopathology dimension ('Anxious-Depression', 'Social-Withdrawal', 'Compulsive-Behaviour and Intrusive Thought') as a within-subjects factor and beta coefficients (averaged across time points) of Instrumental Utility (β1), Hedonic Utility (β2), and Cognitive Utility (β3) as within subject modulating covariates as well as participants' age and gender as between subjects modulating covariates. Once again we observed a significant main effect of Cognitive Utility on psychopathology ($F(1,117) = 4.471$, $p = 0.037$, partial eta square = 0.037). There was no significant effect of Instrumental Utility ($F(1,117) = 1.669$, $p = 0.199$, partial eta square = 0.014) or Hedonic Utility ($F(1,117) = 3.408$, $p = 0.067$, partial eta square = 0.028). No other effects or interactions were significant (all P's > 0.265) except for gender, with females reporting more symptoms (F(2,117) = 4.025, p = 0.020, partial eta square = 0.064). These results suggest that the weight participants' assign to Cognitive Utility, but not the other two utilities, when seeking information is related to their mental health across the three psychopathology dimensions. As in Study 1, doing the analysis on raw numbers does not alter the significance of results.

To illustrate this result in a more simplified manner, we conducted a linear regression with mental health as the dependent measure (quantified as the average psychopathology score across the three dimensions) and the following predictors: the weight assigned to Instrumental Utility (β1) when seeking information, as well as that assigned to Hedonic Utility (β2) and to Cognitive Utility (β3) (all averaged across the two time points). Age and gender were also included as predictors. Confirming the analysis above, a significant inverse relationship was observed between mental health and the weight assigned to Cognitive Utility when seeking information (β = -0.790, p = 0.034), suggesting that participants who seek information more on issues they think of often are the ones who report less psychopathology symptoms across the board. Gender was also significant with females scoring higher on psychopathology symptoms (Gender: β = 0.345, p = 0.005). No other factor was significant (Instrumental Utility: β = 0.498, p = 0.200; Hedonic Utility: β = 0.637, p = 0.063; Age: β = -0.010, p = 0.196; **Figure 2.4**). Finally, correlating each beta with the average psychopathology score across participants (controlling for participants' age and gender), again reveals a significant association with the weight assigned to cognitive utility when seeking information (r(120) = -0.241, p = 0.008), but not with the weight assigned to instrumental (r(120) = 0.114, p = 0.212) or hedonic (trend: r(120) = 0.175, p = 0.053) utilities.

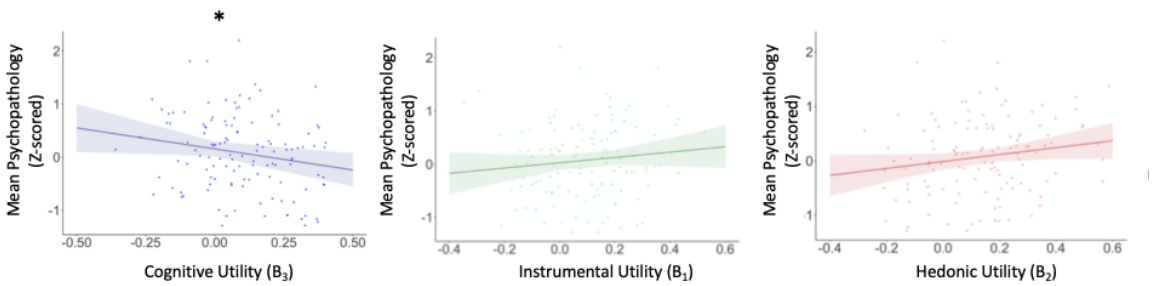

Figure 2.4. Association between information-seeking and mental health is robust to replication. Plotted on the y-axis is the average psychopathology scores across the three dimensions, Z-scored. On the x-axis are the weights assigned to each information-seeking motive from a linear regression predicting information-seeking from Instrumental Utility (green), Hedonic Utility (red) and Cognitive Utility (blue), averaged across the two time points. Dots represent individual subjects. Shading represents confidence interval. Line represents the relationship between the abscissa and ordinate controlling for the effect of the other two motives as well as of age and gender. As can be observed, participants who placed a large positive weight on Cognitive Utility when seeking information reported less psychopathology symptoms (p = 0.034, two-sided), while we observed no effect of Instrumental Utility (p = 0.200,

two-sided) or Hedonic Utility ($p = 0.063$, two-sided). Error bars SEM. *$P < 0.05$ (two-sided). N = 124 subjects.

Across domains information-Seeking is best explained by taking into account Instrumental, Hedonic and Cognitive Utilities (Study 3). We next asked whether the three motives identified in studies 1 and 2 are significantly related to information-seeking in different domains. To that end we conducted a third study in which participants were asked whether they wanted financial information. As in Study 2, this study was longitudinal.

We tested 149 participants on a similar information-seeking task as described above in Study 1 and 2, however here we included 40 stimuli related to finance (e.g., "Do you want to know what the unemployment rate is in Europe?", "Do you want to know the exchange rate between Dollar and Pound?"). Once again, we included all participants who passed the attention check and had enough data variability that allowed us to generate three beta coefficients (that is did not insert the same rating for all stimuli on any of the scales, Time 1 N = 122). Three weeks later, we invited these participants to participate in a follow up study (Time 2). Ninety-five participants completed the follow up study on average 23.43 days following Time 1. Two participants were not included due to providing a different Prolific ID at the two time points. Eighty-two participants passed the attention check and had enough data variability that allowed us to generate three beta coefficients. Descriptive statistics of ratings and their inter-relationships are displayed for Time 1 in **Supplementary Table 2.4** and Time 2 in **Supplementary Table 2.5**.

The data was analysed as in Study 1 and 2. We observed a significant fixed effect of Instrumental Utility (Time 1: $\beta = 0.266 \pm 0.022$ (SE), $t(109.56) = 12.223$, $p < 0.001$, **Figure 2.5a**; Time 2: $\beta = 0.279 \pm 0.029$ (SE), $t(78.98) = 9.497$, $p < 0.001$, **Figure 2.5c**), Hedonic Utility (Time 1: $\beta = 0.094 \pm 0.017$ (SE), $t(106.45) = 5.646$, $p < 0.001$, **Figure 2.5a**; Time 2: $\beta = 0.097 \pm 0.018$ (SE), $t(61.76) = 5.293$, $p < 0.001$, **Figure 2.5c**) and Cognitive Utility (Time 1: $\beta = 0.154 \pm 0.018$ (SE), $t(120.09) = 8.787$, $p < 0.001$, **Figure 2.5a**; Time 2: $\beta = 0.190 \pm 0.022$ (SE), $t(82.98) = 8.473$, $p < 0.001$, **Figure 2.5c**). Once more, the models which included Instrumental Utility, Hedonic Utility and Cognitive Utility as predictors of information-seeking, fit the data better than models including only a subset of the three utilities and also of models including participants' confidence regarding the information to be revealed according to the BIC (**see Figure 2.5b&d**) and the AIC (**see Supplementary Table 2.8**). Note, that at Time 1 two competing models (Hedonic + Instrumental and Hedonic + Cognitive) did not converge.

Once again, most individuals had a dominant motive. 44.26% of individuals assigned more than twice the weight to one motive than the other two motives at Time 1 and 50% at Time 2, and 80.32% of individuals assigning at least 1.25 times more weight to one motive than the other two at Time 1 and 80.49% at Time 2 (**Figure 2.5e&f**). Different motives were dominant for different individuals, with action being dominant for 42.62% of individuals at Time 1, 46.34% at Time 2, affect for 18.03% at Time 1 and 10.98% at Time 1, cognition for 19.67% at Time 1, and 23.17% at Time 2, and 19.67% did not have one particularly strong motive at Time 1 and 19.51% at Time 2 (that is no motive was assigned a weight at least 1.25 greater than the rest).

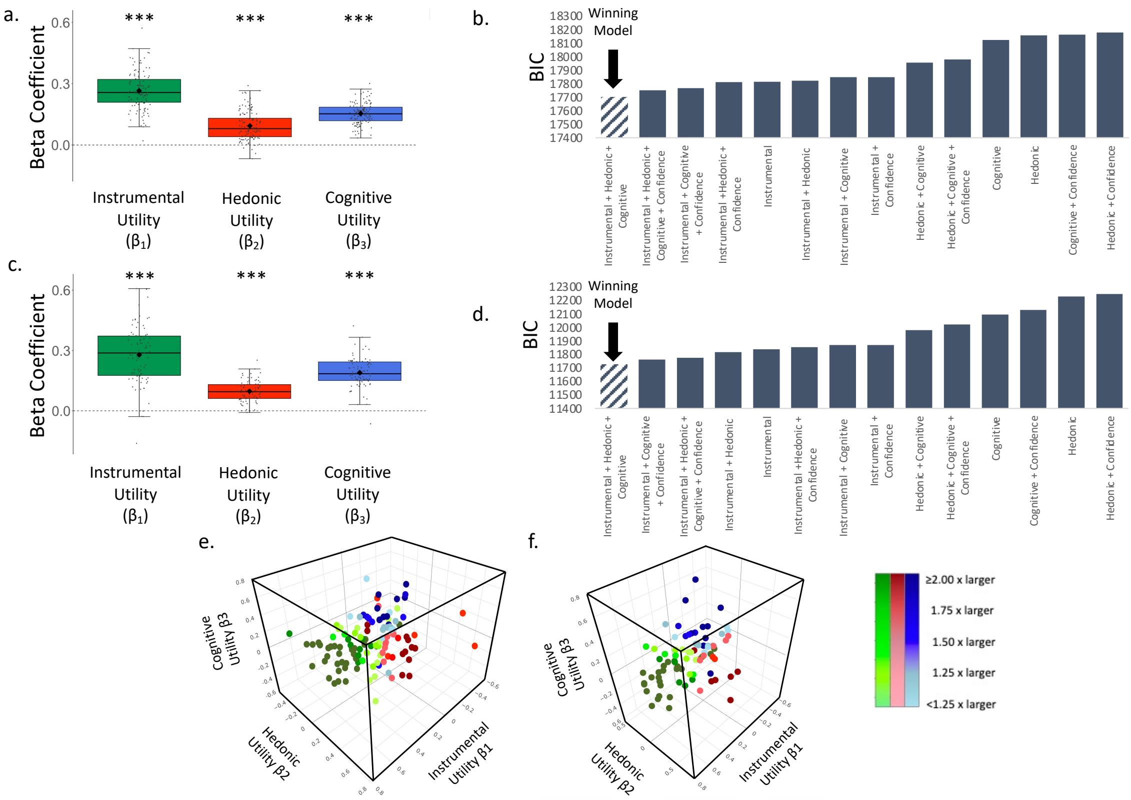

Figure 2.5. Information-Seeking Motives in the Financial Domain. (a&c) Plotted is a boxplot depicting the beta coefficients from a linear mixed effects model at **(a)** Time 1 (N = 122 subjects) and **(c)** Time 2 (N = 82 subjects), which shows that subjects' desire to receive information was greater when the Instrumental Utility (Time 1 $p < 0.001$, Time 2 $p < 0.001$; two-sided), Hedonic Utility (Time 1 $p < 0.001$, Time 2 $p < 0.001$; two-sided) and Cognitive Utility (Time 1 $p < 0.001$, Time 2 $p < 0.001$; two-sided) of information were higher. These were estimated respectively by subjects' ratings of how useful the information would be, how they would feel to know vs not know, and how frequently they think about the stimulus. For each boxplot, the horizontal lines indicate median values, boxes indicate 25–75% interquartile range and whiskers indicate 1.5x interquartile range; individual scores are shown separately as dots. **(b&d)** BIC scores from **(b)** Time 1 and **(d)** Time 2 reveal that the model described in **(a&c)** fit the data better than models including other combinations of the utilities and those including subjects' confidence regarding what the information would reveal. The same was true when examining AIC scores (**see Supplementary Table 2.8**). Smaller BIC and AIC scores indicate better fit. **(e&f)** Plotted are the weights each individual put on each motive when seeking information at **(e)** Time 1 and **(f)** Time 2. Beta coefficients of Instrumental Utility are on the x-axis, of Cognitive Utility on the y-axis and of Hedonic Utility on the z-axis. Green dots represent participants who put the largest weight on Instrumental Utility when seeking information. Red dots represent participants who put the largest weight on Hedonic Utility when seeking information. Blue dots represent

participants who put the largest weight on Cognitive Utility when seeking information. The colour gradient represents how dominant the largest weight was in comparison to the other two weights. Individuals who put more than twice as much weight on their dominant utility than the other two utilities are represented in darkest colours. Those whose dominant utility was less than 1.25 times larger than the other two are represented in the lightest colours. *** = $P < 0.001$ (two-sided).

We next tested to what extent the relative weight of the three information-seeking motives are stable over time within individuals. First, we measured by how much each participant moved over time within the three-dimensional space plotted in **Figures 2.5e&f**. This indicates changes in the relative weights a participant assigned to the three betas. We then tested whether the magnitude of change was significantly smaller than chance. To test this, we reran the exact same analysis above for each subject, but each time mismatching one participant's T1 data with another participant's T2 data (i.e., permutation test). We then compared the average distance participants actually moved in the three-dimensional space from T1 to T2 to the average distance calculated from the permutation test. We did this 10,000 times and found that 100% of the times the average distance participants actually moved from T1 to T2 was smaller than chance (mean difference between iterations and actual mean movement = 0.087, range = 0.015-0.15). Second, we calculated the Intraclass Correlation Coefficient (ICC) of each beta type across the time points (see methods). The ICC across time was significant for Instrumental Utility (ICC = 0.317, $p = 0.044$) and Hedonic (ICC = 0.329, $p = 0.039$) utilities, but not for Cognitive Utility (ICC = 0.019, $p = 0.446$), suggesting that the weights assigned to frequency of thought, while stable across time in the self-trait domain, is not in the finance domain. The weight assigned to expected affect and instrumental utility when seeking information show some stability across time in both the financial and social domains.

Stability of information-seeking motives across domains (Study 4). Next, we wanted to know whether the three motives identified in studies 1-3 significantly predicted information-seeking in a third domain, health, and whether these motives were stable within an individual across domains. To investigate this, we conducted a fourth study in which we invited 101 new participants as well as all participants who completed Study 3, Time 1 (N = 122) to complete another information-seeking task, but this time in the domain of Health. One-hundred and forty-eight participants completed the study, which included 47 participants from Study 3, Time 1 (which was conducted on average 166 days previous).

The task was similar to Study 1-3, however here we included 40 stimuli related to health (e.g., "Would you like to know if you have a gene that increases your likelihood of Alzheimer's disease?", "Would you like to know if you have a gene that increases your likelihood of a Strong Immune System?"). Once again, data was analysed for all participants who passed the attention check and who had enough data variability that allowed us to generate three beta coefficients (that is did not insert the same rating for all stimuli on any of the scales, N = 116). Descriptive statistics of ratings and their inter-relationships are displayed in **Supplementary Table 2.6**.

The data was analysed as in Study 1, 2 and 3. We observed a significant fixed effect of Instrumental Utility ($\beta = 0.229 \pm 0.026$(SE), $t(126.52) = 8.918$, $p < 0.001$, **Figure 2.6a**), Hedonic Utility ($\beta = 0.090 \pm 0.020$ (SE), $t(103.74) = 4.447$, $p < 0.001$, **Figure 2.6a**) and Cognitive Utility ($\beta = 0.096 \pm 0.015$ (SE), $t(128.60) = 6.295$, $p < 0.001$, **Figure 2.6a**). Once more, the models which included Instrumental Utility, Hedonic Utility and Cognitive Utility as predictors of information-seeking, fit the data better than models including only a subset of the three utilities and also of models including participants' confidence regarding the information to be revealed according to the BIC (**see Figure 2.6b**) and the AIC (**see Supplementary Table 2.8**).

Once again, most individuals had a dominant motive. 52.59% of individuals assigned more than twice the weight to one motive than the other two motives, while 89% of individuals assigned at least 1.25 times more weight to one motive than the other two (**Figure 2.6c**). Different motives were dominant for different individuals, with action being dominant for 57.76% of individuals, affect for 19.83%, cognition for 11.21% and 11.21% did not have one particularly strong motive (that is no motive was assigned a weight at least 1.25 greater than the rest).

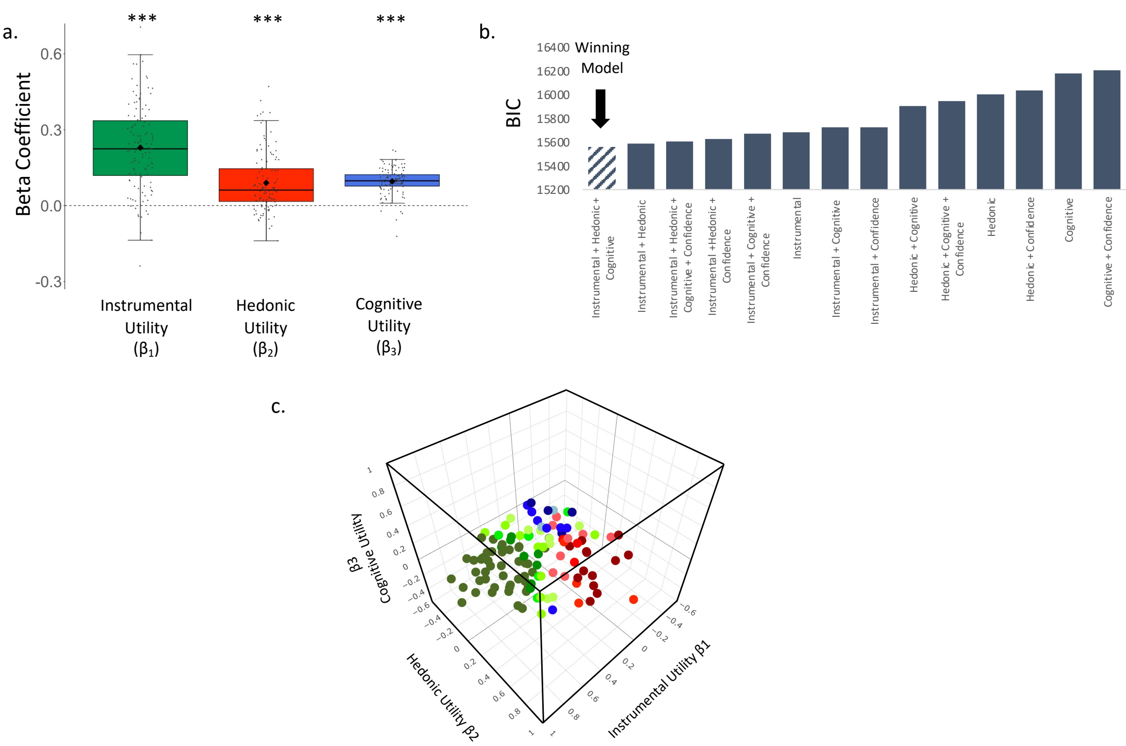

Figure 2.6. Information-Seeking Motives in the Health Domain. (a) Plotted are the beta coefficients from a linear mixed effects model (two-sided; N = 116 subjects), showing that subjects' desire to receive health related information was greater when the Instrumental Utility ($p < 0.001$, two-sided), Hedonic Utility ($p < 0.001$, two-sided) and Cognitive Utility ($p < 0.001$, two-sided) of information were higher. These were estimated respectively by subjects' ratings of how useful the information would be, how they would feel to know vs not to know, and how frequently they think about the stimulus. The horizontal lines indicate median values, boxes indicate 25–75%

interquartile range and whiskers indicate 1.5 × interquartile range; individual scores are shown as dots. **(b)** BIC scores reveal that the model described in **(a)** fit the data better than models including alternate combinations of the utilities and also those including subjects' confidence regarding what the information would reveal. The same was true when examining AIC scores (**see Supplementary Table 2.8**). Smaller BIC and AIC scores indicate better fit. **(c)** Plotted are the weights each individual put on each motive when seeking information in the health domain. Beta coefficients of Instrumental Utility are on the x-axis, of Cognitive Utility on the y-axis and of Hedonic Utility on the z-axis. Green dots represent participants who put the largest weight on Instrumental Utility when seeking information. Red dots represent participants who put the largest weight on Hedonic Utility when seeking information. Blue dots represent participants who put the largest weight on Cognitive Utility when seeking information. The colour gradient represents how dominant the largest weight was in comparison to the other two weights. Individuals who put more than twice as much weight on their dominant utility than the other two utilities are represented in darkest colours. Those whose dominant utility was less than 1.25 times larger than the other two are represented in the lightest colours. *** = $P < 0.001$, ** = $P < 0.01$ (two-sided).

We next tested to what extent the relative weight of the three information-seeking motives are stable across domain (i.e., Finance and Health) and time within individuals. Data was analysed for all those participants who completed Study 3, Time 1 and Study 4 and who passed the attention check and who had enough data variability that allowed us to generate three beta coefficients (N = 38). We first measured by how much they moved over domain/time within the three-dimensional space plotted in **Figures 2.5e & 2.6c**. This indicated changes in the relative weights a participant assigned to the three betas. We then tested whether the magnitude of change was significantly smaller than chance. To test this, we reran the exact same analysis above for each subject, but each time mismatching one participant's Study 3, Time 1 data with another participant's Study 4's data (i.e., permutation test). We then compared the average distance participants actually moved in the three-dimensional space from Study 3, Time 1 to Study 4 to the average distance calculated from the permutation test. We did this 10,000 times and found that 99.73% of the times the average distance participants actually moved from Study 3, Time 1 to Study 4 was smaller than chance (mean difference between iterations and actual mean movement = 0.08, range = 0.02-0.17). Second, we calculated the Intraclass Correlation Coefficient (ICC) of each beta type across the time points (**see Methods**). As in Study 3, the ICC for Instrumental Utility (ICC = 0.621, p = 0.002) and Hedonic Utility (ICC = 0.445, p = 0.042) was significant, and the ICC for Cognitive Utility (ICC = 0.272, p = 0.172) was not. Note, on average, the action motive was greater in health and finance domains than the self-trait domain. These findings together indicated that while 'trait' impacts the importance people assign to information-seeking motives other factors such as state and domain may matter too.

2.4 Discussion

The desire for knowledge is a fundamental part of human nature (Kidd & Hayden, 2015). People spend a substantial amount of time actively pursuing information, for example by asking questions, reading or conducting online searches. These activities, often referred to as 'information-seeking' behaviours, are integral to learning, social engagement and decision-making (Loewenstein, 1994; Gottlieb et al., 2013; Sakaki et al., 2018).

Here we show that people want information more when they believe information (i) will be useful in guiding their actions, (ii) will have a positive impact on their affective state and (iii) is related to concepts they often think about. A model which incorporates these three motives, reflecting the influence participants expect information to have on their action, affect and cognition, explained individuals' information-seeking choices better than a range of other models. These results were replicated across four studies and three different domains – information about self-traits, finance, and health – suggesting that the model is likely domain general.

We observed individual differences with regard to the weights participants assigned to the three motives when seeking information. Many participants assigned a particularly large weight to one of the motives relative to the other two. That is, some participants were driven mostly to seek information according to its (i) predicted usefulness (action-driven), (ii) its predicted impact on their feelings (affect-driven), (iii) while others mostly sought information that relate to concepts they think of frequently (cognitive-driven). The individual differences in the weight people assign to the different motives when seeking information can help explain individual differences in what people want to know (Sharot & Sunstein, 2020). For example, a subject who assigns more weight to instrumental utility than hedonic utility may be more inclined to want to know if they have a predisposition to breast cancer than a subject who assigns more weight to hedonic utility than instrumental utility.

Our longitudinal studies indicate that these individual differences are fairly stable over time. Moreover, in Study 4 (which included a smaller sample size than the other studies) we found that individuals who tended to assign a large weight to a motive in one domain (i.e., finance) relative to other individuals, tended to do so in another domain (e.g., health). We also saw interesting differences across domains, with the action motive being much greater in health and finance domains than the self-trait domain, and the cognitive motive being much more stable across time within the self-trait domain and not finance domain. Together, these findings suggest that the weights people assign to the different motives are likely a combined function of trait, state and domain.

The individual differences in the weights participants assigned to the three motives were related to mental health within the domain of self-traits. Specifically, those individuals who assigned a larger weight to the cognitive motive when seeking information reported less psychopathology symptoms across the board. The theory suggests that this relationship should hold true in other domains as well (i.e., finance and health), given the relative stability of these motive weights across these domains.

However, this aspect remains to be empirically tested. Our approach differs from the past few attempts to test for a relationship between information-seeking and psychopathology (Aderka et al., 2013; Hildebrand-Saints & Weary, 1989; Camp, 1986; Locander & Hermann, 1979; Gray & Tonge, 200; Giancardo et al., 2016) in two fundamental ways. First, rather than examining for an association between psychopathology and the frequency of information-seeking (an approach which has led to mixed results; Aderka et al., 2013; Hildebrand-Saints & Weary, 1989; Camp, 1986; Locander & Hermann, 1979; Gray & Tonge, 2001), we examined for an association between mental health and participants' motives for seeking information. Second, instead of using traditional psychopathology diagnosis, we adopted a dimensionality approach (Gillan et al., 2016; Rouault et al., 2018). This approach considers the possibility that a specific symptom is predictive of several psychiatric conditions, allowing an investigation that cuts through classic clinical boundaries. Our results suggest that the relative importance of the information-seeking motives about the self are related to general mental health.

We have previously theorised that the relationship between mental health and information seeking is bidirectional (Sharot and Sunstein, 2020). Our study, however, is correlational and thus we cannot conclude whether certain patterns of information-seeking lead to increase/decrease in psychopathology symptoms, and/or the other way around. Moreover, our findings suggest that the three motives measured here are associated with information-seeking but cannot speak of causation. We also note that according to our theory (Sharot & Sunstein, 2020) people first predict what information will likely reveal and based on that prediction estimate utilities. In some situations, expected information can be quantified and is highly correlated with a quantifiable estimated utility. For example, a person's expectations on how they will be rated on intelligence by others will be correlated with how they expect to feel when they receive that information (i.e., if they expect to be rated positively, they will probably feel good knowing the rating). In this specific case, a researcher could interchangeably use expected information or expected affect to predict information seeking choice, because the former is simply the subjective assessment of the latter. In most cases, however, the two are not easily interchangeable. For example, if a person expects information will reveal the Dollar to Pound exchange rate is high that on its own does not tell us how they likely expect to feel about such information.

I recognise that assessing what participants' 'want to know' within the confines of a controlled lab experiment might not completely reflect their information-seeking choices in real-world scenarios where consequences to such choices exist. As a next step, it's important to explore these behaviours through naturalistic experiments, where participants' motivations for seeking information can be observed in settings that more closely resemble their everyday environments. This approach would help bridge the gap between controlled experimental findings and the nuanced realities of real-life information-seeking behaviour.

In sum, we have provided evidence that people's decisions about whether to seek or avoid information are related to an integration of the instrumental value, hedonic value, and cognitive value of information. We further show that individual differences in information-seeking reflect varying emphasis on these values, which in turn provides clues about participants mental health. These findings could be used to facilitate policy makers' ability to calculate the costs and benefits of information disclosure (Sunstein, 2018; Thunström, 2019). Moreover, it suggests that by presenting information in a way that taps into the three motives of information-seeking, policy makers may increase the likelihood that individuals will engage with and benefit from vital information.

2.5 References

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.

Afifi, W. A., & Weiner, J. L. (2004). Toward a theory of motivated information management. Communication Theory, 14(2), 167-190.

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.

Allport, G. W., & Odbert, H. S. (1936). Trait-names: A psycho-lexical study. Psychological monographs, 47(1), i.

Anderson, T. H., & Schwartz, R. M. (1992). Intolerance of ambiguity and depression: Evidence of their relationship. Journal of Abnormal Psychology, 101(2), 278-281.

Barr, D. J., Levy, R., Scheepers, C., & Tily, H. J. (2013). Random effects structure for confirmatory hypothesis testing: Keep it maximal. Journal of memory and language, 68(3), 255-278.

Berlyne, D. E. (1957). Uncertainty and conflict: a point of contact between information-theory and behavior-theory concepts. Psychological Review, 64(6p1), 329.

Bozdogan, H. (1987). Model selection and Akaike's information criterion (AIC): The general theory and its analytical extensions. Psychometrika, 52(3), 345-370.

Camp, C. J. (1986). I am curious-grey: Information seeking and depression across the adult lifespan. Educational Gerontology, 12(4), 375-384.

Caspi, A., Houts, R. M., Belsky, D. W., Goldman-Mellor, S. J., Harrington, H., Israel, S., ... & Moffitt, T. E. (2014). The p factor: one general psychopathology factor in the structure of psychiatric disorders?. *Clinical psychological science*, *2*(2), 119-137.

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, 115(31), E7255-E7264.

Christensen, P. N., Stein, M. B., & Means-Christensen, A. (2003). Social anxiety and interpersonal perception: A social relations model analysis. Behaviour Research and Therapy, 41(11), 1355-1371.

Cuthbert, B. N., & Insel, T. R. (2010). Toward new approaches to psychotic disorders: the NIMH Research Domain Criteria project.

Cuthbert, B. N., & Insel, T. R. (2013). Toward the future of psychiatric diagnosis: the seven pillars of RDoC. BMC medicine, 11(1), 1-8.

Doron, G., Moulding, R., Kyrios, M., & Nedeljkovic, M. (2008). Sensitivity of self-beliefs in obsessive compulsive disorder. Depression and Anxiety, 25(10), 874-884.

Eager, C., & Roy, J. (2017). Mixed effects models are sometimes terrible. arXiv preprint arXiv:1701.04858.

Foa, E. B., Liebowitz, M. R., Kozak, M. J., Davies, S., Campeas, R., Franklin, M. E.,... & Tu, X. (2005). Randomized, placebo-controlled trial of exposure and ritual prevention, clomipramine, and their combination in the treatment of obsessive-compulsive disorder. American Journal of psychiatry, 162(1), 151-161.

Fresco, D. M., Coles, M. E., Heimberg, R. G., Liebowitz, M. R., Hami, S., Stein, M. B., & Goetz, D. (2001). The Liebowitz Social Anxiety Scale: a comparison of the psychometric properties of self-report and clinician-administered formats. Psychological medicine, 31(6), 1025-1035.

Garner, D. M., Olmsted, M. P., Bohr, Y., & Garfinkel, P. E. (1982). The eating attitudes test: psychometric features and clinical correlates. Psychological medicine, 12(4), 871-878.

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(1), 1-10.

Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A., & Daw, N. D. (2016). Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. Elife, 5, e11305.

Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information avoidance. Journal of Economic Literature, 55(1), 96-135.

Gottlieb, J., Oudeyer, P. Y., Lopes, M., & Baranes, A. (2013). Information-seeking, curiosity, and attention: computational and neural mechanisms. Trends in cognitive sciences, 17(11), 585-593.

Graham, K. T., Martin-Iverson, M. T., Holmes, N. P., Jablensky, A., & Waters, F. (2014). Deficits in agency in schizophrenia, and additional deficits in body image, body schema, and internal timing, in passivity symptoms. Frontiers in Psychiatry, 5, 126.

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.

Hards, E., Ellis, J., Fisk, J., & Reynolds, S. (2020). Negative view of the self and symptoms of depression in adolescents. Journal of affective disorders, 262, 143-148.

Hertwig, R., & Engel, C. (2016). Homo ignorans: Deliberately choosing not to know. Perspectives on Psychological Science, 11(3), 359-372.

Hildebrand-Saints, L., & Weary, G. (1989). Depression and social information gathering. Personality and Social Psychology Bulletin, 15(2), 150-160.

Hirshleifer, J., & Riley, J. G. (1979). The analytics of uncertainty and information-an expository survey. Journal of economic literature, 17(4), 1375-1421.

Jacobi, C., Paul, T., de Zwaan, M., Nutzinger, D. O., & Dahme, B. (2004). Specificity of self-concept disturbances in eating disorders. International Journal of Eating Disorders, 35(2), 204-210.

Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. Journal of personality and social psychology, 103(1), 54.

Kappes, A., Harvey, A. H., Lohrenz, T., Montague, P. R., & Sharot, T. (2020). Confirmation bias in the utilization of others' opinion strength. Nature neuroscience, 23(1), 130-137.

Karlsson, N., Loewenstein, G., & Seppi, D. (2009). The ostrich effect: Selective attention to information. Journal of Risk and uncertainty, 38(2), 95-115.

Kidd, C., & Hayden, B. Y. (2015). The psychology and neuroscience of curiosity. Neuron, 88(3), 449-460.

Klayman, J., & Ha, Y. W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. Psychological review, 94(2), 211.

Kobayashi, K., & Hsu, M. (2019). Common neural code for reward and information value. Proceedings of the National Academy of Sciences, 116(26), 13061-13066.

Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. Nature human behaviour, 3(6), 587-595.

Kreps, D. M., & Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. Econometrica: journal of the Econometric Society, 185-200.

Lerman, C., Hughes, C., Lemon, S. J., Main, D., Snyder, C., Durham, C., ... & Lynch, H. T. (1998).'What you don't know can hurt you: adverse psychologic effects in members of BRCA1-linked and BRCA2-linked families who decline genetic testing. Journal of Clinical Oncology, 16(5), 1650-1654.

Locander, W. B., & Hermann, P. W. (1979). The effect of self-confidence and anxiety on information seeking in consumer risk reduction.

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. Psychological bulletin, 116(1), 75.

Marin, R. S., Biedrzycki, R. C., & Firinciogullari, S. (1991). Reliability and validity of the Apathy Evaluation Scale. Psychiatry research, 38(2), 143-162.

Mason, O., Linney, Y., & Claridge, G. (2005). Short scales for measuring schizotypy. Schizophrenia research, 78(2-3), 293-296.

Matuschek, H., Kliegl, R., Vasishth, S., Baayen, H., & Bates, D. (2017). Balancing Type I error and power in linear mixed models. Journal of memory and language, 94, 305-315.

Murayama, K., Sakaki, M., Yan, V. X., & Smith, G. M. (2014). Type I error inflation in the traditional by-participant analysis to metamemory accuracy: a generalized mixed-effects model perspective. Journal of Experimental Psychology: Learning, Memory, and Cognition, 40(5), 1287.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. Review of general psychology, 2(2), 175-220.

Patton, J. H., Stanford, M. S., & Barratt, E. S. (1995). Factor structure of the Barratt impulsiveness scale. Journal of clinical psychology, 51(6), 768-774.

Persoskie, A., Ferrer, R. A., & Klein, W. M. (2014). Association of cancer worry and perceived risk with doctor avoidance: an analysis of information avoidance in a nationally representative US sample. Journal of behavioral medicine, 37(5), 977-987.

Rouault, M., Seow, T., Gillan, C. M., & Fleming, S. M. (2018). Psychiatric symptom dimensions are associated with dissociable shifts in metacognition but not task performance. Biological psychiatry, 84(6), 443-451.

Sakaki, M., Yagi, A., & Murayama, K. (2018). Curiosity in old age: A possible key to achieving adaptive aging. Neuroscience & Biobehavioral Reviews, 88, 106-116.

Sass, L. A., & Parnas, J. (2003). Schizophrenia, consciousness, and the self. Schizophrenia bulletin, 29(3), 427-444.

Saunders, J. B., Aasland, O. G., Babor, T. F., De La Fuente, J. R., & Grant, M. (1993). Development of the alcohol use disorders identification test (AUDIT): WHO collaborative project on early detection of persons with harmful alcohol consumption-II. Addiction, 88(6), 791-804.

Schwarz, G. J. T. A. O. S. (1978). The annals of statistics. Estimating the dimension of a model, 6, 461-464.

Seow, T. X., & Gillan, C. M. (2020). Transdiagnostic phenotyping reveals a host of metacognitive deficits implicated in compulsivity. Scientific reports, 10(1), 1-11.

Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. Nature Human Behaviour, 4(1), 14-19.

Shukla, P., & Solanki, R. (2013). Web based keystroke dynamics application for identifying emotional state. Int. J. Adv. Res. Comput. Commun. Eng, 2(11), 4489-4493.

Silverstone, P. H., & Salsali, M. (2003). Low self-esteem and psychiatric patients: Part I–The relationship between low self-esteem and psychiatric diagnosis. Annals of general hospital psychiatry, 2(1), 1-9.

Spielberger, C. D., & Gorsuch, R. L. (1983). State-trait anxiety inventory for adults: Manual and sample: Manual, instrument and scoring guide. Consulting Psychologists Press
Stigler, G. J. (1961). The Economics of Information. p. Journal of Political Economy, 69.
Sunstein, C. R. (2018). The cost-benefit revolution. MIT Press.
Sunstein, C. R. (2019). Ruining popcorn? The welfare effects of information. Journal of Risk and Uncertainty, 58(2), 121-142.
Thunström, L. (2019). Welfare effects of nudges: The emotional tax of 57calorie menu labeling. Judgment and Decision making, 14(1), 11.
Trudel, N., Scholl, J., Klein-Flügge, M. C., Fouragnan, E., Tankelevitch, L., Wittmann, M. K., & Rushworth, M. F. (2021). Polarity of uncertainty representation during exploration and exploitation in ventromedial prefrontal cortex. Nature Human Behaviour, 5(1), 83-98.
van Lieshout, L. L., Vandenbroucke, A. R., Müller, N. C., Cools, R., & de Lange, F. P. (2018). Induction and relief of curiosity elicit parietal and frontal activity. Journal of Neuroscience, 38(10), 2579-2588.
Zung, W. W. (1965). A self-rating depression scale. Archives of general psychiatry, 12(1), 63-70.
Yoshie, M., & Haggard, P. (2013). Negative emotional outcomes attenuate sense of agency over voluntary actions. Current Biology, 23(20), 2028-2032.
Owens, K. M., Asmundson, G. J., Hadjistavropoulos, T., & Owens, T. J. (2004). Attentional bias toward illness threat in individuals with elevated health anxiety. Cognitive Therapy and Research, 28, 57-66.

Chapter 3

Information-Seeking Reflects and Shapes Well-Being

3.1 Overview

Determining which factors are associated with well-being has been a key pursuit of scientists, policymakers, and the general public. Research has linked well-being to various elements such as social relationships (Pieh et al., 2020; Ertel et al., 2009; Robles & Kiecolt-Glaser, 2003), exercise (Marconcin et al., 2022; Peluso et al., 2005), and wealth (Ettman et al., 2022; Pollack et al., 2007). In recent years, as people spend more time online, the need to investigate the relationship between online activity and well-being has become imperative (DataReportal, 2022). This will inform the development of online tools to enhance well-being and could provide real-time assessment of it (e.g., digital phenotyping).

One of the most frequent online activities is information-seeking. Interestingly, what people choose to know varies vastly from one individual to the next (Kelly & Sharot, 2021; Kobayashi et al., 2019; Sunstein, 2019). These variations may provide important clues about an individual's inner cognitive and affective state (Kelly & Sharot, 2021). In particular, we have theorised that the affective properties of the knowledge people consume from self-guided searches may reflect and shape their well-being (Sharot & Sunstein, 2020). In other words, the relationship between well-being and information-seeking may be reciprocal and form a self-reinforcing loop.

Regarding the first direction of this hypothesis—one's affective state influencing the affective characteristics of information sought online—this hypothesis aligns with findings that people with depression tend to engage with stimuli that perpetuate their sadness (Milgram et al., 2015). This mechanism is analogous to that hypothesised to underlie rumination (Watkins, 2008; Michl et al., 2013). In addition, this hypothesis is consistent with rich literature showing that affect alters information seeking and decision making (Sharot & Sunstein, 2020; Kelly & Sharot et al., 2021; Lerner et al., 2015; Hockey et al., 2000; George et al., 2016; Paulus & Angela, 2012; Phelps et al., 2014; Pictet et al., 2011; Charpentier et al., 2018).

In the second direction—how the affective characteristics of browsed information impact well-being—previous research has shown a robust relationship between exposure to negative words and well-being (Mathews & MacLeod, 2005; Lyubomirsky et al., 1998). For example, Lyubomirsky and colleagues (1998) had participants read negative or positive passages and those who read negative passages reported lower mood levels afterward. Thus, it is reasonable to think that affective characteristics of words on webpages that people expose themselves to will have an impact on their well-being.

In a digital age where online activity is increasingly shaping our experiences and perceptions, understanding its impact on our mental and emotional well-being is not just an academic pursuit but a societal imperative. If indeed a bi-directional relationship exists between the type of information that is consumed from self-guided online searches and well-being, it would have significant theoretical and practical implications. As humans constantly engage in information-seeking online, there is a unique opportunity to harness this data to detect mental health issues and help guide information-seeking patterns. Given this rich potential, it is surprising how limited our knowledge is of the links between well-being and the affective properties of information browsed online.

We define well-being based on the American Psychological Association (APA) dictionary, which describes it as *"a state of happiness and contentment, with low levels of distress, overall good physical and mental health and outlook, or good quality of life"* (VandenBos, 2007). In this study, we focus on psychological and emotional subjective well-being, which we assess through self-reported questionnaires evaluating mental health and/or mood. Over four studies (total N = 947) we test the hypothesis that the affective characteristics of the information people expose themselves to online reflect and shape their psychological and emotional subjective well-being. To quantify the affective properties of the information people expose themselves to, we asked participants to share their web-browsing history, ensuring their privacy, and then used a natural language processing (NLP) approach to quantify the valence of the text on webpages that participants browsed. We first related these affective characteristics to participants' emotional and psychological well-being (Study 1 and 2) and then manipulated these factors to examine for causal relationships (Study 3). Finally, we examined whether providing cues about the potential emotional impact of webpages on well-being would influence participants' web-browsing behaviour, in a way that was consistent with improvements in well-being (Study 4).

3.2 Methods

Study 1

Participants. Three hundred and twelve participants completed a study online via Prolific's recruitment platform. Data from 23 participants whose searches did not result in at least 1KB of text from at least 3 webpages each day was not analysed further. Thus, data of 289 participants were analysed (age = 33.17, SD =11.71; females = 50.5%, males = 48.1%, other = 1.4%). Out of those, 171 participants also completed state mood ratings. Data of five participants who indicated that contrary to the instructions they submitted archived browsing history was not included in mood analysis, as their current mood ratings obviously could not be temporally associated with their submitted browsing data, leaving for mood analysis N = 164 (age = 33.23, SD =11.62; females = 52.4%, males = 47.6%, other = 0%)All participants received £7.50 for their participation on day 1 and £3.25 for days 2-5. Ethical approval was provided by the Research Ethics Committee at University College London.

Procedure

Data collection. Participants were asked to browse the internet for 20-30 minutes a day for 5 days using Mozilla Firefox and then submit their internet search history for this period (**see Figure 3.1**). We extracted the paragraph text from each webpage, denoted by <p> in the webpage's html code, using the '*rvest*' package in RStudio. We then cleaned the text by removing extraneous information such as punctuation, symbols (e.g., @, #), emojis, links (URLs), and all other non-alphanumeric characters (similar to Kelley & Gillan, 2022). Participants were asked to browse the internet during non-work hours so that their web-browsing behaviour would not reflect mandatory work-related tasks. All consecutive duplicate webpages were removed from analysis. Participants for whom we had less than three webpages from which we could extract at least 1KB of data per day were excluded from analysis.

Text valence analysis. To quantify the valence of webpages visited, we used the NRC VAD lexicon (Mohammad, 2018), which categorises the valence of terms on a scale from 0 (most negative) to 1 (most positive). In line with Kiritchenko and colleagues (2020), we computed the percentage of words with a Positive valence score greater or equal to 0.75 (2668 terms, e.g., '*delicious*' and '*admire*') and percentage of words with a Negative valence score less or equal to 0.25 (3081 terms, e.g., '*despise*' and '*danger*'), out of all words contained in the extracted text of each webpage visited for each of the five days. We then averaged these positive valence and negative valence scores separately across all webpages visited on each day and then averaged the daily scores across the five days to create a Positive Valence score and Negative Valence score, respectively. We also quantified separately the percentage of *Anger, Fear, Anticipation, Trust, Surprise, Sadness, Joy* and *Disgust* associated words greater or equal to 0.75, as defined by the NRC Emotion Lexicon (Mohammad and Turney, 2013), out of all words on each webpage visited by participants for each day and then across days (i.e., Emotion scores).

To assess whether the NRC VAD lexicon scores of webpages was related to alternative sentiment analysis approaches, we computed the Intraclass Correlation Coefficient (ICC) for the mean NRC Positive and Negative Valence scores of webpages visited by participants (N = 100) separately with the same from another widely used lexicon, the Hu and Liu Opinion lexicon (Hu & Liu, 2004), which categorises 2006 words as positive and 4783 words as negative. Next, we calculated the ICC for the mean NRC Positive and Negative scores with a state-of-the-art large language machine learning model, the distilbert-base-uncased-finetuned-sst-2-english (i.e., Distilbert; HuggingFace, 2022). Finally, we calculated the ICC for the Hu and Liu Opinion lexicon with the distilbert-base-uncased-finetuned-sst-2-english.

To examine if the NRC VAD lexicon scores corresponded to human ratings, we asked participants (N = 100) to rate the positive (0 (not at all) to 6 (very positive)) and negative valence (0 (not at all) to 6 (negative)) of 10 randomly assigned webpages from a corpus of 48 webpages. We then computed the NRC Positive and Negative Valence score for each webpage and their respective human rating for that webpage and submitted the positive and negative pairs into an ICC to calculate their reliability.

Next, we were interested whether the valence of webpages' whole text was associated with the valence of a sample of its text. To test this, we randomly extracted segments from webpages (N = 100) with a minimum word count of 200 words. We then calculated the Positive and Negative scores for the random samples text and that of its corresponding whole text and submitted those into an ICC analysis to calculate reliability (separately for the Positive and Negative scores). We also examined if there is good reliability between the valence scores of the text of a whole webpage and the valence of the text that participants attended to the most. To test this, a new group of participants were asked to browse the internet for 10-minutes, while their eye movements were tracked via a web camera (**see** https://app.gazerecorder.com). We tested ten participants who in total browsed 31 websites. We quantified the NRC valence scores for the text participants attended to the most on each webpage (e.g., indicated as red in the heatmap by the algorithm) and the scores of the whole text on the webpage and submitted those into an ICC analysis to calculate reliability (separately for the Positive and Negative scores).

Mental Health and mood. On day one, participants completed self-report questionnaires which assess psychopathology symptoms (the list is adopted from Gillan et al., 2016) These included: Obsessive-Compulsive Inventory – Revised (OCI-R; Foa et al., 2005), Self-Rating Depression Scale (SDS; Zung, 1965), State-Trait Anxiety Inventory (STAI; Spielberger et al., 1983), Alcohol Use Disorder Identification Test (AUDIT; Saunders et al., 1993), Apathy Evaluation Scale (AES; Marin et al., 1991), Eating Attitudes Test (EAT-26; Garner et al., 1982), Barratt Impulsivity Scale (BIS-11; Patton et al., 1995), Short Scales for Measuring Schizotypy (Mason et al., 2005), Liebowitz Social Anxiety Scale (LSAS; Fresco et al., 2001). On days 1-5, participants indicated their current mood directly before their web-browsing session and directly afterwards, on scales from -50 (very unhappy) to + 50 (very happy). The task was coded using the Qualtrics online platform (https://www.qualtrics.com).

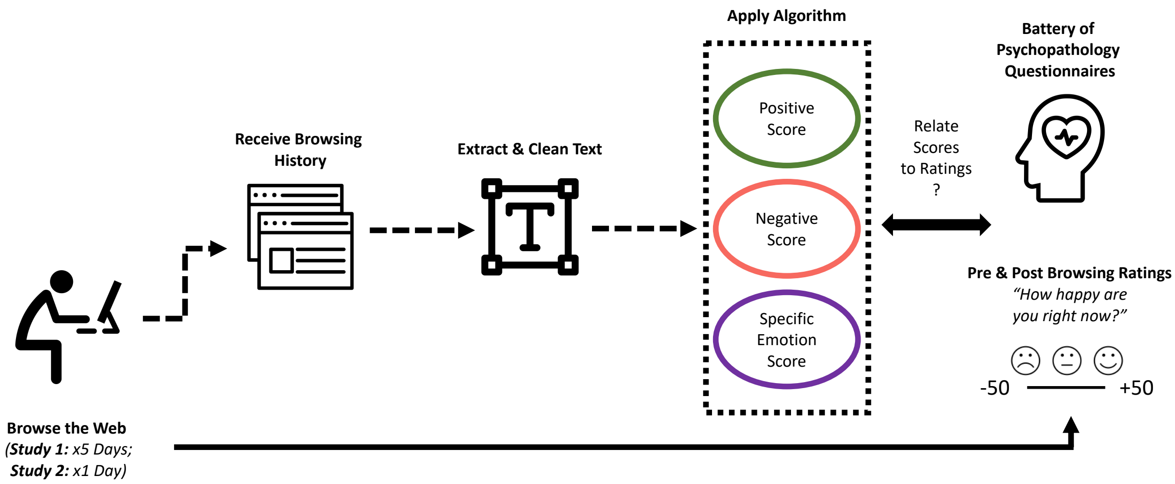

Figure 3.1. Data collection and pre-processing pipeline. Participants browsed the internet for 20-30 minutes a day for 1-5 days using Mozilla Firefox and then submitted their internet search history for this period. We extracted the paragraph text from each webpage, denoted by < p > in the webpage's html code and cleaned it (**see Methods**).

The text was then submitted to an algorithm that calculated a Negative score and a Positive score for each webpage (**see Methods**) as well as scores for Anger, Fear, Anticipation, Trust, Surprise, Sadness, Joy and Disgust (**see Supplementary Chapter 3**). On day one participants completed self-report questionnaires which assess mental health. On days 1-5, participants also indicated their mood directly before and after the web-browsing session. Participates' scores were then related to self-reported psychopathology symptoms and mood.

Analysis

Relating the valence of webpages to mental health. Each participant was scored on the three psychopathology dimensions identified by Gillan and colleagues (2016) and replicated by Rouault and colleagues (2018) ('*Anxious-Depression*', '*Social-Withdrawal*' and '*Compulsive-Behaviour and Intrusive Thought*'. To generate these scores, we followed Kelly and Sharot (2021) - we first Z-scored the ratings for each questionnaire item separately across participants. Next, we multiplied each Z-scored item by its factor weight as identified earlier (Gillan et al., 2016). Then for each subject the three-psychopathology dimension scores were calculated by summing all of the weighted items assigned to each dimension.

The valence of webpages visited by participants were then related to the psychopathology dimensions by submitting the three-psychopathology dimension scores into a mixed ANOVA with psychopathology dimension as a within-subject factor and valence as within subject modulating covariates as well as participants' age and gender as between-subjects modulating covariates (similar to Kelly & Sharot, 2021). This analysis was then followed up with a simplified analysis in which the average of the three-psychopathology dimension scores of each individual were entered as a dependent measure in a linear regression with valence entered as an independent measure as well as age and gender.

Relating the valence of webpages to mood. To investigate the relationship between web-browsing patterns and mood, we asked participants to indicate their current mood directly before their web-browsing session and directly afterwards, on scales from -50 ('*very unhappy*') to +50 ('*very happy*'). We first assessed whether participants pre-browsing mood was related to the valence of information they browsed. To that end, we ran two separate mixed effect models each including participants pre-browsing mood ratings as fixed and random effects along with age and gender as fixed effect predicting the Negative Valence score and Positive Valence score of webpages visited, separately. Next, we were interested in whether the valence of the webpages that participants browsed had an impact on their mood directly after browsing the internet. To test this, we once again ran two mixed effect models, each predicting post browsing mood ratings from either the Negative and Positive Valence score of webpages visited (input as a fixed and random effect), controlling for pre-browsing mood (fixed and random effect) as well as age and gender (fixed effect).

Assessing the stability of the valence of web-browsing across time. To assess the within-subject stability of the valence of webpages visited across the five days, we calculated an intraclass correlation coefficient (ICC). Specifically, we submitted separately the Negative and Positive Valence score and the scores for the specific emotions of webpages visited by each participant for each of the five days into ICC analysis.

Study 2: Replication of Study 1

Participants. Five hundred participants completed a study online via Prolific's online recruitment system. Data of 53 participants from whom we could not obtain at least 1KB of text from a minimum of 3 webpages a day was not analysed. Thus, data of 447 participants were analysed (age = 33.85, SD =12.58; females = 56.4%, males = 41.8%, other = 1.8%). For the mood analysis, we only included those participants that submitted data that was browsed during the study session (N = 400, age = 33.23, SD =11.62; females = 52.4%, males = 47.6%, other = 0%), as otherwise their reported mood ratings would not be temporally reflective of their submitted browsing data. Participants received £7.50 for their participation. Ethical approval was provided by the Research Ethics Committee at University College London.

Procedure

The procedure was exactly as in Study 1 except that all participants were asked to browse the internet for 30-minutes for one day.

Analysis

Relating the valence score of webpages to psychopathology. This analysis was conducted as described in Study 1.

Relating the valence of webpages to mood. We first tested whether participants pre-browsing mood was related to the valence of information they browsed. As we only had one observation per participant for each variable of interest, (compared to five observations in Study 1), we ran two simple linear regressions predicting the Negative Valence score and Positive Valence score, separately, from pre-browsing mood ratings, controlling for age and gender. Next, we were interested in whether the valence of the webpages that participants browsed had an impact on their mood directly after browsing the internet. To test this, we ran two simple linear regressions, both predicting participants post browsing mood ratings from either the Negative or Positive Valence score of webpages visited. Both models controlled for participants pre-browsing mood ratings, age and gender.

Study 3

Participants. One hundred and thirty-nine participants completed the study on Qualtrics (www.qualtrics.com) and were recruited via Prolific's online recruitment platform (www.prolific.co). Participants received £7.50 per hour for their participation. Thirty-seven participants were excluded for not providing at least 3 webpages from which we could extract at least 1KB of data, leaving 102 participants (negative valence condition: N = 55, age = 33.96, SD = 9.68; females = 45.5%, males = 49.1%, other =

5.5%; control condition: N = 47, age = 34.72, SD = 12.14; females = 46.8%, males = 51.1%, other = 2.1%).

Procedure

Data collection. To assess the directionality of the relationship between mood and web-browsing patterns, we first conducted a manipulation of webpages that participants were exposed to. Specifically, we asked participants to browse two webpages, randomly selected from either six very negative (i.e., negative valence manipulation) or six positive webpages. The stimuli were selected from webpages that participants browsed in studies 1-2. The valence of the webpages was quantified using the exact method as outlined in Study 1 and pages were included if they had a Negative score greater or equal to 2.5 standard deviations from the mean (i.e., negative webpages) or a Negative score between 0 and 1 standard deviations from the mean (i.e., neutral pages). Participants indicated their happiness levels on a scale ranging from very unhappy (-50) to very happy (+50) before and after the manipulation.

Next, participants were asked to browse the internet for 10-minutes using Mozilla Firefox and then submit their internet search history for this period. We then extracted the paragraph text from each webpage, denoted by <p> in the webpage's html code, using the '*rvest*' package in RStudio. All consecutive duplicate webpages were removed from analysis.

Analysis

To assess whether the mood manipulation was successful, a 2x2 ANOVA with condition (negative manipulation, control) as a between-subject factor, and time (pre-manipulation, post manipulation) as the within-subject factor was conducted. Follow up pair-wise t-tests were also conducted. Next, for each participant, we computed the Negative Valence score of the webpages browsed. Finally, we tested for a difference in the Negative Valence score of the webpages browsed between the negative valence manipulation group and control group.

<u>Study 4</u>

Participants. One hundred and nine participants (label condition: N = 55; no label condition: N = 54) completed the study on Qualtrics (www.qualtrics.com) and were recruited via Prolific's online recruitment platform (www.prolific.co). Participants received £7.50 per hour for their participation.

Procedure

Data collection. Participants were assigned to either a label condition or no label condition. In the **no label condition** participants were randomly presented with three Google search result pages from a set of 18. Each page contained three possible webpage links they could click on. They simply had to click on one of the three on each trial. They would then spend 90 seconds browsing that webpage.

Participants in the **label-condition** did the same, except that next to each link there was a label indicating the affective label of that webpage. The label was assigned based on valence scores calculated as in studies 1-2. If the Positive score of the page was >2.5 SD from the mean of webpages browsed in studies 1-2, the webpage was given the label *'feel better'*; If the negative score of the page was >2.5 SD from the mean of webpages browsed in studies 1-2, the webpage was given the label *'feel worse';* if neither was true it was given the label neutral. The labels indicate that *on average* this website makes people feel worse/better.

Analysis

To assess whether the manipulation was successful, we used a 2 x 3 ANOVA with condition (label vs. no label) as a between-subject factor, and label valence (positive, negative, neutral) as the within-subject factor. Follow up pair-wise t-tests were conducted.

3.3 Results

Below we report the observed associations between the affective properties of information browsed online and subjective well-being (Study 1 and 2) and then manipulate the factors of interest to test for causation (Study 3). Finally using insight from Study 1, 2 and 3, we develop an intervention to alter information-seeking patterns (Study 4). Note, that Study 2 is a replication of Study 1 except that Study 2 includes one observation per participant compared to five observations in Study 1.

Information-Seeking is Associated with Well-Being (Study 1 & 2). Participants in Study 1 (N = 289) browsed the web for 20-30 minutes a day for five days, and in Study 2 (N = 447) for 30 minutes on one day. They then submitted their web-browsing history. We used this web browsing history to access the web pages visited and extracted the text of these websites (**see Methods**). We then scored the text on affective properties (positive and negative valence, and specific emotions; **see Figure 3.1**).

Quantifying the affective properties of web pages. There are many validated methods to score text on sentiment (valence). These include machine-learning methods (Devlin et al., 2008; Goldberg et al., 2014; Liu et al., 2019) and 'bag of words' (lexicon) approaches which are developed by asking large groups of people to rate words on specific dimensions (Hu & Liu, 2004; Mohammad, 2018). We first tested whether these different methods provide consistent scores for participants. We selected two popular lexicons - the NRC VAD lexicon (Mohammad, 2018) and the Hu and Liu Opinion lexicon (Hu & Liu, 2004) - and a state-of-the-art large language machine learning model, the distilbert-base-uncased-finetuned-sst-2-english (i.e., Distilbert; HuggingFace, 2022; **see Methods** for details). We used each method separately to score all webpages visited by the first 100 participants from Study 1 and averaged the webpage scores for each participant. We used an intra-class correlation coefficient (ICC) analysis to examine how consistent the scores were across different scoring methods, separately for positive and negative scores. We observed good reliability between all three methods: (i) the NRC VAD lexicon and the Hu and Liu Opinion lexicon (Positive score: ICC = 0.835, $p < 0.001$**;** Negative score: ICC = 0.948, $p < 0.001$); (ii) the NRC VAD lexicon and the Distilbert algorithm (Positive score: ICC = 0.812, $p < 0.001$**;** Negative score: ICC = 0.869, $p < 0.001$); and (iii) the Distilbert algorithm and the Hu and Liu Opinion lexicon (Positive score: ICC = 0.866, $p < 0.001$**;** Negative score: ICC = 0.885, $p < 0.001$). This suggests that these different methods measure the same construct.

As the NRC lexicons performed equivalently to machine learning algorithms but required significantly less computational resources, we opted to use it. We first checked that the valence ratings computed by the NRC VAD lexicon were reflective of human assessment. One hundred participants, each rated 10 webpages on how positive and negative they were. These scores were significantly related to the NRC Valence scores (Negative score: ICC = 0.707, $p < 0.001$, Positive score: ICC = 0.499, $p < 0.001$), suggesting that the NRC VAD Lexicon scores reflect human subjective assessment of webpages well.

Given that the method we used scores entire webpages rather than the sub-text participants consume, it was important to test whether the former was a good indicator of the latter. To that end we adopted two approaches. First, we examined whether there is good reliability between the valence of text on a whole webpage and the valence of text on random part of it. To test this, we randomly extracted segments of text from webpages (N = 100) with a minimum word count of 200 words (see Yazman, 2017). We observed good reliability between the NRC Valence scores of randomly sampled segments and the scores of their respective webpage's whole texts (Negative score: ICC = 0.968, $p < 0.001$; Positive score: ICC = 0.908, $p < 0.001$). This result suggests that by analysing the whole text of a webpage, we can reliably compute the sentiment of a random section of a webpage. Second, we examined directly if there is good reliability between the valence scores of the text of a whole webpage and the valence of the text that participants attended to the most. To test this, a new group of participants were asked to browse the internet for 10-minutes, while their eye movements were tracked via a web camera (**see** https://app.gazerecorder.com). We tested ten participants who in total browsed 31 websites. We quantified the NRC Valence scores for the text that participants attended to the most on each webpage (i.e., indicated as red in the heatmap by the algorithm) and the scores of the whole text on the webpage. There was good reliability between the NRC Valence scores of the text attended to the most and the text of the whole webpage (Positive score: ICC = 0.832, $p < 0.001$; Negative score: ICC = 0.760, $p < 0.001$, **Figure 3.2b-c**). Thus, the valence scores of the whole text of a webpage is a good indicator of the valence of text that participants attend to the most. Together these checks suggest that using the NRC VAD lexicon is a suitable measure to assess the valence of information on webpages that participants consume.

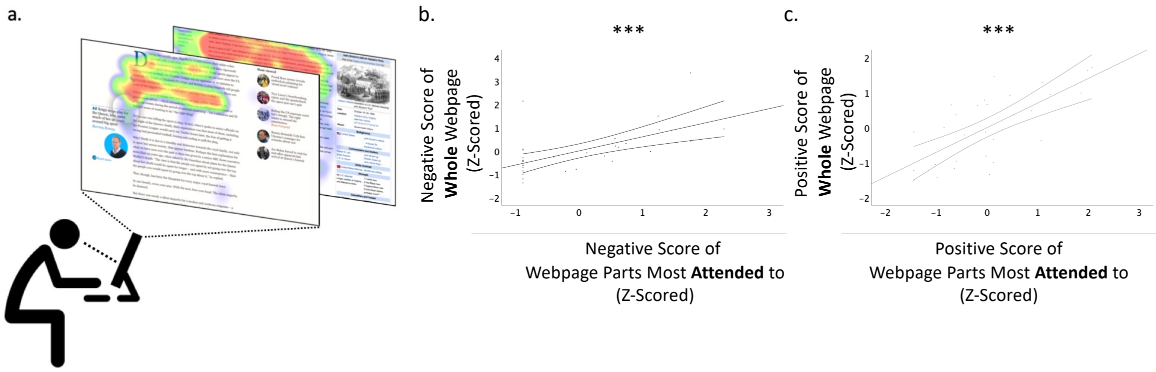

Figure 3.2. The valence of the whole text of webpages is reflective of the valence of the text participants attend to the most. (a) Participants (N = 10) were asked to browse the internet for 10-minutes, while their eye movements were tracked via a web camera (**see https://app.gazerecorder.com**). For both **(b)** Negative scores and **(c)** Positive scores, there was strong reliability between the NRC scores of the text participants attended to the most on the website (e.g., indicated as red in the heatmap by the algorithm) (y-axis) and the score of the whole text of a website (x-axis; black). Thus, computing the valence of the whole text of webpages is a good indicator of the valence of text that participants consumed. Outer lines represent confidence intervals. Inner line represents the relationship between the abscissa and ordinate. Each dot represents a webpage. *** = P < 0.001 (two-sided).

Quantifying mental health. To assess mental health, we adopted a dimensionality approach, which considers the possibility that a specific psychopathology symptom is predictive of several conditions and allows an investigation that cuts through classic clinical psychopathology boundaries (Gillan et al., 2016; Rouault et al., 2018; Seow & Gillan, 2020). In particular, previous work used a factor analysis across items in a large battery of traditional psychopathology questionnaires and identified three psychopathology dimensions across those items: '*Anxious-Depression*', '*Social-Withdrawal*' and '*Compulsive-Behaviour and Intrusive Thought'* (Gillan et al., 2016). The factor analysis provided a weight for each item in relation to each dimension. Thus, a person's symptom severity for each dimension can be quantified by having an individual complete a battery of traditional psychopathology questionnaires and then calculating a weighted average across items' ratings. Indeed, this is what we did for each participant; we Z-scored the ratings of each questionnaire item separately across participants and then for each participant we calculated the three-dimension scores as explained above (as done in Kelly & Sharot, 2021, **see Methods** for more details).

Affective properties of webpages visited provides a marker of mental health. We first examined whether the tendency to browse content with a specific valence was stable over time. To that end, we calculated the Intraclass Correlation Coefficient (ICC) of the Negative and Positive Valence of webpages visited by each participant over the five days. The ICC of both the Negative score (ICC = 0.554, p < 0.001) and Positive score (ICC = 0.626, p < 0.001) all indicate moderate stability, which was statistically significant. This suggests that the tendency is likely impacted both by 'trait-like' and 'state-like' tendencies.

We next examined if there is a relationship between mental health and the affective properties of pages participants browsed. For each participant, we calculated the three-psychopathology dimension scores ('*Anxious-Depression*', '*Social-Withdrawal*', '*Compulsive-Behaviour and Intrusive Thought*'), which we submitted to a within-subjects factors mixed ANOVAs. In the first mixed ANOVA, the Negative Valence score of the webpages that participants browsed (Z-scored) was input as a within-subject modulating factor. Participants' age and gender were entered as between-subject modulating covariates (both Z-scored). We observed a significant main effect of the Negative Valence score of webpages participants browsed on psychopathology scores (**Study 1:** $F(1,284) = 4.464$, $p = 0.035$, partial eta square = 0.015; **Study 2:** $F(1,442) = 8.303$, $p = 0.004$, partial eta square = 0.018). These results suggest that individuals who browse webpages that are more negatively valenced, experience poorer mental health across the three mental health dimensions. The Negative Valence of webpages browsed is thus a general fingerprint of mental health, rather than associated with a specific condition.

To show this result in a more intuitive manner, we conducted a linear regression with psychopathology as the dependent measure (quantified as the average psychopathology score across the three dimensions) and the Negative Valence score of the webpages that participants browsed, age and gender as the predictor variables, all Z-scored. In line with the results above, we observed a significant positive relationship between psychopathology and the Negative Valence score of webpages participants browsed (**Study 1:** $\beta = 0.087 \pm 0.042$ (SE), $t(288)= 2.069$, $p = 0.039$, $r = 0.122$, **Figure 3.3a**; **Study 2:** $\beta = 0.099 \pm 0.034$ (SE), $t(446) = 2.930$, $p = 0.004$, $r = 0.138$, **Figure 3.3b**), suggesting that participants who browsed more negatively valenced webpages reported worse mental health.

The second mixed ANOVA was identical to the first except that the Positive Valence score was input as a within-subject modulating factor instead of the Negative Valence score. We did not observe a significant main effect of Positive Valence score of webpages on psychopathology scores in Study 1, (**Study 1:** $F(1,284) = 0.000$, $p = 0.997$, partial eta square = 0.000), although there was a significant effect in Study 2 (**Study 2:** $F(1,442) = 8.149$, $p = 0.005$, partial eta square = 0.018) with participants reporting higher psychopathology symptoms browsing less positively valanced text.

We implemented the exact same method described above using a second valence lexicon (Hu & Liu, 2004). All results were replicated (**see Supplementary Analysis, Chapter 3**), suggesting that the results are not restricted to a specific method.

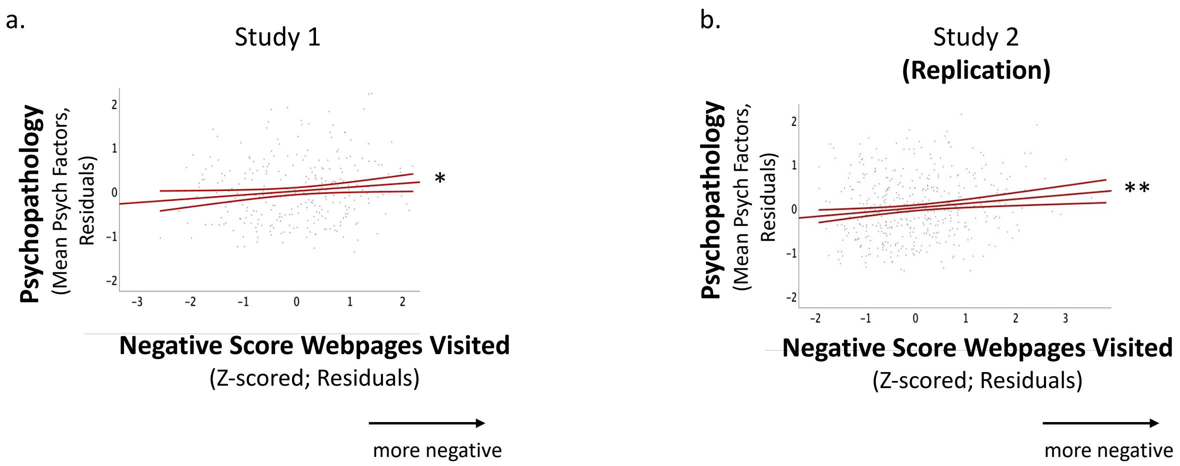

Figure 3.3. Self-guided browsing of negative content online is associated with poorer mental health. (a&b) Greater psychopathology symptoms (the average score across the three dimensions) are associated with higher Negative Valence score in **(a)** Study 1 and **(b)** Study 2. Dots represent the residual values from the model for individual participants. The outer lines represent confidence intervals. The inner line represents the relationship between the abscissa and ordinate controlling for the effect of age and gender. **P < 0.01, *P < 0.05 (two-sided).

As participants knew they would submit their browsing history, it is possible they may browse differently than if 'no one was watching', despite anonymity. This would induce noise that may make the relationship between information-seeking and well-being more difficult to detect and thus likely even larger than reported here. While participants were explicitly asked to browse the internet during the study session, not all of them did. We suspected as much from some of the participants short study completion times and thus asked participants after completing the study, whether they indeed submitted data that was browsed in-session or from their archived browsing history. Thirty-nine participants in Study 1 and 7 in Study 2 admitted they submitted archived data (due to this small N the following analysis was conducted across Study 1 and 2 together). We tested whether the average Valence scores of webpages browsed from this group was different than for those who browsed in-session – it was not for Positive scores (browsed in study data M = 0.015, SD = 0.032, archived data M = 0.015, SD = 0.030, t(514) = 0.243, p = 0.808, Cohen's d = 0.107) nor for Negative scores (browsed in study data M = 0.030, SD = 0.014, archived data M = 030, SD = 0.014, t(514) = 0.112, p = 0.928, Cohen's d = 0.081). This suggests that participants were not browsing more positive or negative webpages on average due to the study set-up.

A Bidirectional Association between Information-Seeking and Mood (Study 1 & 2). Thus far, we observed that the valence of information consumed from self-guided searches provides a general fingerprint of mental health. Next, we ask whether it is also associated with mood, which is a feature of well-being, and if so whether this association is bidirectional.

To that end, we asked participants (**Study 1**: N = 164; **Study 2**: N = 400) to indicate their current mood directly before their web-browsing session and directly afterwards, on scales from -50 (very unhappy) to +50 (very happy). We tested whether participants pre-browsing mood and post-browsing mood was related to the valence of information they browsed.

First, we examined the relationship between valence of information consumed and pre-browsing mood. We ran separate linear mixed effect models - one predicting the Negative Valence score and the other the Positive Valence score of webpages visited - from participants' pre-browsing happiness ratings in Study 1 (fixed and random effects) along with age and gender as fixed effects. In Study 2, as we only had one observation per subject for each variable of interest (compared to five in Study 1), we ran two simple linear regressions predicting the Negative Valence score and Positive Valence score from pre-browsing mood ratings, controlling for age and gender (**see Supplementary Table 3.1** for control variable statistics). We found that participants who reported better mood prior to browsing the internet, exposed themselves to less negatively valenced webpages (**Study 1:** β -0.082 $\pm$ 0.041 (SE), $t(380.29) = -1.981$, $p = 0.048$, **Figure 3.4a**; **Study 2:** β = -0.096 $\pm$ 0.049 (SE), $t(399) = -1.974$, $p = 0.049$, $r = -0.099$, **Figure 3.4a**), with no significant relationship observed for Positive Valence score (**Study 1:** $\beta = -0.001 \pm 0.002$ (SE), $t(104.88) = -0.493$, $p = 0.623$; **Study 2:** $\beta = 0.088 \pm 0.048$ (SE), $t(399) = 1.830$, $p = 0.068$, $r = 0.092$).

Next, we ran a similar analysis as above to predict post-browsing mood from the Negative score of webpages participants visited, while controlling for pre-browsing mood in all models. We found that participants expressed better mood after browsing less negatively valenced webpages, controlling for mood pre-browsing age and gender (**Study 1:** $\beta = -0.044 \pm 0.019$ (SE), $t(58.12) = -2.338$, $p = 0.023$, **Figure 3.4b**; **Study 2:** $\beta = -0.093 \pm 0.035$ (SE), $t(399) = -2.686$, $p = 0.008$, $r = -0.134$, **Figure 3.4b**). Participants also reported better mood after browsing more Positive Valence webpages in Study 1 (**Study 1:** $\beta = 0.037 \pm 0.019$ (SE), $t(82.62) = 2.013$, $p = 0.047$) but this effect was not significant in Study 2 ($\beta = 0.063 \pm 0.035$ (SE), $t(399) = 1.770$, $p = 0.077$, $r = 0.089$). Together, these results suggest a bi-directional relationship between mood and the Negative Valence of webpages participants consume from self-guided se arches. Specifically, individuals that were happier directly before browsing the internet, browsed less negatively valenced information, and individuals who browsed less negatively valenced information reported being happier after browsing the internet. As these results are still correlational, we next ran a study to test for causation (**see Study 3**).

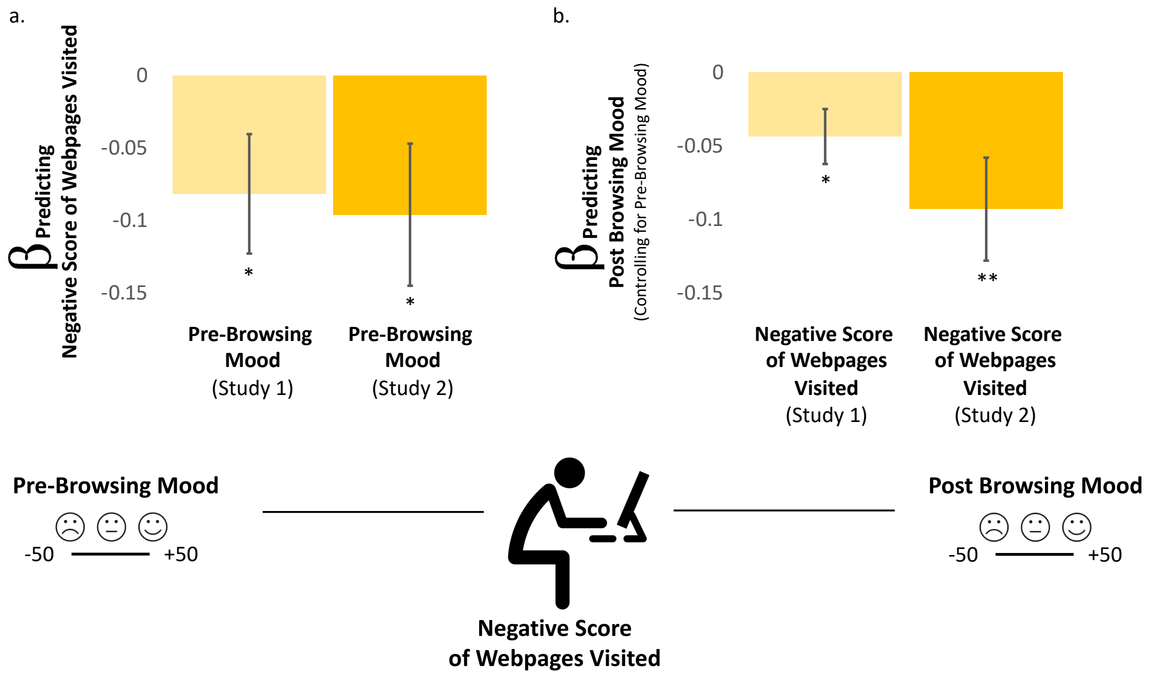

Figure 3.4. Browsing more negatively valanced webpages is associated with worse mood before and after browsing. (a) Plotted on the y-axis is the beta coefficient predicting the Negative score of webpages visited by participants from their pre-browsing mood in Study 1 (light yellow) and Study 2 (dark yellow). Participants with worse pre-browsing mood tend to browse more negatively valenced webpages controlling for age and gender in both studies. **(b)** Plotted on the y-axis is the beta coefficient predicting participants post browsing mood from the Negative score of webpages they visited in Study 1 (light yellow) and Study 2 (dark yellow), controlling for pre-browsing mood, age and gender. Participants who browsed more negatively valenced webpages reported worse post-browsing mood. Error bars = standard error (SEM). ** = $P < 0.01$, * = $P < 0.05$ (two-sided).

The Bidirectional Association between Information-Seeking and Mood is Causal (Study 3). In Study 3 (N = 102), we tested whether the relationship between browsing negatively valenced information and well-being is causal. To do so, we first manipulated the webpages participants were exposed to and then tested for mood. Specifically, participants were asked to read information from two webpages randomly selected either from six negative webpages (i.e., negative valence condition; N = 55) or six neutral pages (i.e., control condition; N = 47). The negative pages were randomly selected from all webpages browsed in Study 1 that were +2.5 SD from the mean Negative score. The neutral webpages were randomly selected from webpages browsed in Study 1 that were between -1 and +1 SD from the mean. Participants indicated their mood levels on a scale ranging from '*very unhappy'* (-50) to '*very happy'* (+50) before and after being exposed to the webpages.

A 2 (condition: negative valence, control) by 2 (time: pre-manipulation, post manipulation) ANOVA on self-reported mood revealed a significant interaction ($F(1, 97) = 15.922$, $p < 0.001$, partial eta squared = 0.141). The interaction was characterised by participants in the negative valence condition reporting feeling unhappier post manipulation ($M = -1.93$, $SD = 23.69$) compared to pre-manipulation ($M = 9.47$, $SD = 24.29$, $t(54) = -5.031$, $p < 0.001$, Cohen's $d = -0.678$), with no difference in the control condition (post manipulation: $M = 9.31$, $SD = 19.31$, pre-manipulation: $M = 9.53$, $SD = 19.42$, $t(46) = 0.131$, $p = 0.896$, Cohen's $d = 0.019$). Importantly, participants in the negative valence condition reported feeling unhappier post manipulation relative to controls (negative valence condition: $M = -1.93$, $SD = 23.69$; control condition: $M = 9.53$, $SD = 19.43$, $t(100) = 2.242$, $p = 0.010$, Cohen's $d = 0.525$), with no difference pre-manipulation (negative valence condition: $M = 9.47$, $SD = 24.29$; control condition: $M = 9.32$, $SD = 19.32$, $t(100) = -0.035$, $p = 0.972$, Cohen's $d = -0.007$; **see Figure 3.5a**). This suggests that being exposed to negatively valence webpages results in worse mood.

Now that the negative valence group reported worse mood than the control group, we asked whether this group of participants would go on to consume more negatively valenced webpages than the control group from self-guided searches. To that end, participants were asked to browse the internet for 10-minutes and then submit their internet search history for this period. The negative valence of webpages participants exposed themselves to was quantified as in studies 1 and 2 (**see Methods**). Results show that participants in the negative valence condition subsequently browsed significantly more negatively valenced webpages ($M = 0.034$, $SD = 0.020$) than those in the control condition ($M = 0.026$, $SD = 0.014$, $t(96.04) = -2.259$, $p = 0.026$; Cohen's $d = -0.436$; **see Figure 3.5b**). These results suggest a causal bi-directional relationship between participants' mood and web-browsing patterns (**see Figure 3.5c**). All results remain the same when removing participants that have a values plus/minus 3 standard deviations from the mean.

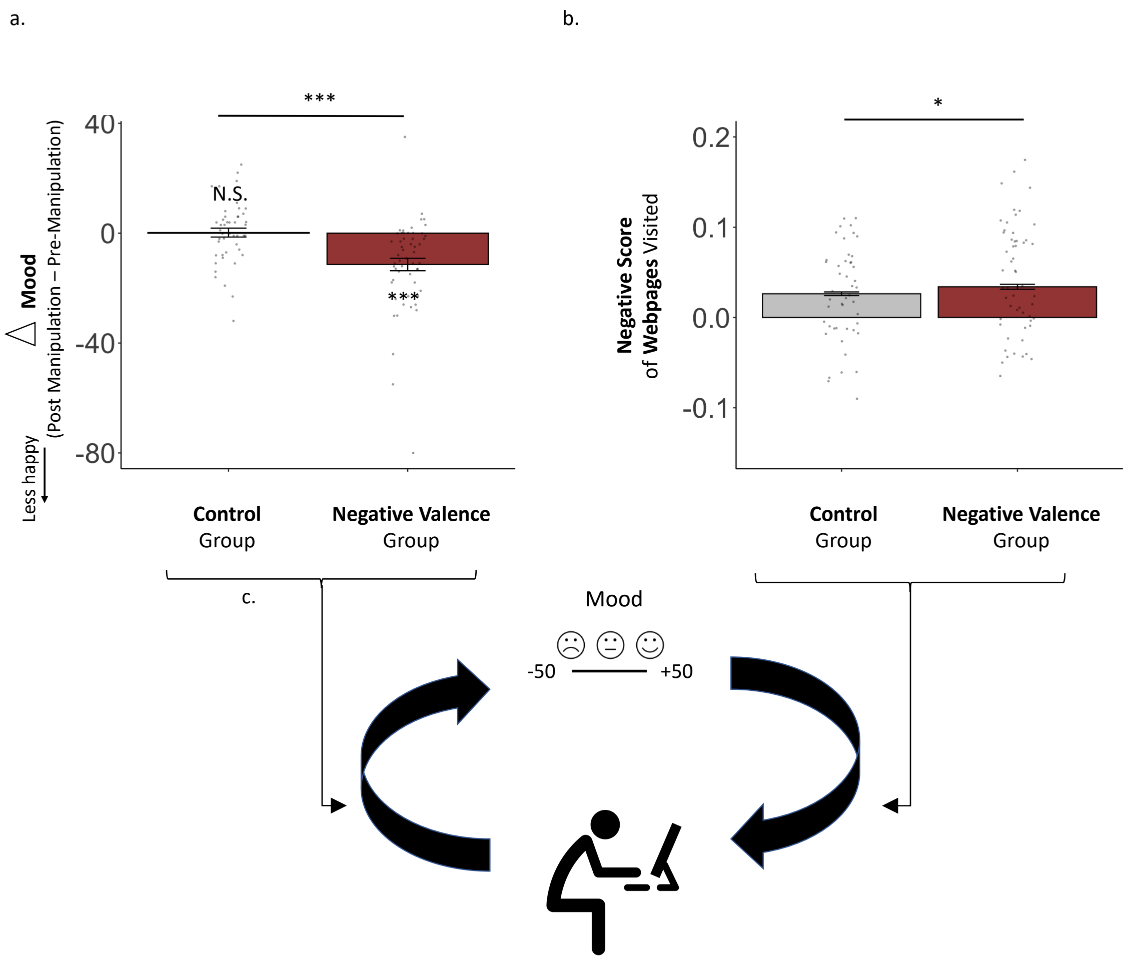

Figure 3.5. Bi-directional relationship between mood and the valence of information consumed. (a) Participants were asked to browse two webpages, randomly selected from either six very negative webpages or six neutral webpages (control). Participants reported their mood on a scale ranging from 'very unhappy' (-50) to 'very happy' (+50) before and after the manipulation. Plotted on the y-axis are participants' post manipulation mood rating minus their pre-manipulation mood rating for the negative valence condition (grey) and control condition (red). Participants in the negative valence group reported worse mood after browsing compared to before, while participants in the control condition reported no difference in their mood after browsing compared to before. Moreover, participants in the negative valence condition reported worse mood after browsing than those in the neutral condition. **(b)** After browsing the webpages selected by us participants had the opportunity to freely browse the web. Those in the negative valence condition browsed significantly more negatively valenced webpages than those in the control condition. Individual scores are shown as dots. **(c)** The results suggest a bi-directional relationship between mood and valence of webpages browsed. Specifically, **(b)** worse mood leads to browsing more negatively valenced information, and **(a)** browsing more negatively valenced information leads to worse mood. Error bars = standard error (SEM). *** = P < 0.001, * = P < 0.05, N.S. = not significant (two-sided).

An Intervention to Alter Patterns of Information-Seeking (Study 4). Studies 1-3 show that browsing negatively valanced information is associated with negative features of psychological and emotional well-being. We thus pondered whether people would select to expose themselves to more positive and less negative information if they had advance knowledge of the affective properties of webpages. That is, would providing people with cues about the valence of webpages alter their information-seeking patterns, resulting in less consumption of negative information and more consumption of positive information?

To answer this question, we conducted Study 4. Participants were assigned to either a label condition or no label condition. In the **no label condition** participants were randomly presented with three Google search result pages from a set of 18 (**Figure 3.6a**). Each page contained three possible webpage links participants could click on. They simply had to click on one of the three on each trial. They would then spend at least 90 seconds browsing that webpage. These 18 pages were selected from Google's list of frequent queries, for which Google results contained varying levels of valence scores (i.e., positive, neutral and negative).

Participants in the **label condition** did the same, except that next to each link there was a label indicating the sentiment of that webpage (**Figure 3.6a**). The label was assigned based on valence scores calculated as in studies 1-2. If the Positive score of the page was >2.5 SD from the mean of webpages browsed in studies 1-2, the webpage was given the label 'feel better'; If the negative score of the page was >2.5 SD from the mean of webpages browsed in studies 1-2, the webpage was given the label 'feel worse'; if neither was neither it was given the label 'neutral'. The labels indicate whether on average this website makes people feel worse/better.

The question of interest was if participants would use the labels to alter the information, they exposed themselves to. The results suggest they did. A 2 (condition: label, no label) by 3 (valence: positive, neutral, negative) ANOVA on webpage choices revealed a significant interaction between condition and valence ($F(1, 107) = 7.695$, $p = 0.007$, partial eta squared = 0.067; **see Figure 3.6b**). The interaction was characterised by participants in the label condition selecting more webpages with the positive label ($M = 1.444$, $SD = 1.04$) than the no label condition ($M = 1.055$, $SD = 0.68$, $t(90.93) = -2.314$, $p = 0.023$, Cohen's $d = -0.445$; **see Figure 3.6b**) and less webpages with the negative label ($M = 0.630$, $SD = 0.73$) than the no label condition ($M = 1.000$, $SD = 0.839$, $t(107) = 2.251$, $p = 0.016$, Cohen's $d = 0.469$; **see Figure 3.6b**). There was no difference in the number of neutral webpages selected between the label condition ($M = 0.910$, $SD = 0.88$) and no label condition ($M = 0.910$, $SD = 0.85$, $t(107) = 0.010$, $p = 0.992$, Cohen's $d = 0.002$; **see Figure 3.6b**).

Additionally, within the label condition webpages with the positive label were selected more than neutral webpages (Mean Difference = 0.537, $SD = 1.77$, $t(53) = 2.234$, $p = 0.030$, Cohen's $d = 0.304$) and negative label webpages (Mean Difference = 0.815, $SD = 1.58$, $t(53) = 3.792$, $p < 0.001$, Cohen's $d = 0.516$), with the latter two not different (Mean Difference = 0.278, $SD = 1.23$, ($t(53) = 1.653$, $p = 0.104$, Cohen's $d = 0.225$).

In contrast, in the no label condition none of the webpages were labelled, thus there was no difference in the likelihood of selecting webpages which should have been labelled as positive and neutral (Mean Difference = 0.145, SD = 1.28, $t(54)$ = 0.841, p = 0.404, Cohen's d = 0.113), or should have been labelled as positive and negative (Mean Difference = 0.055, SD = 1.27, $t(54)$ = 0.319, p = 0.751, Cohen's d = 0.043), nor between those which should have been labelled as neutral and negative (Mean Difference = -0.091, SD = 1.53, $t(54)$ = -0.440, p = 0.661, Cohen's d = -0.059).

Together, the results suggest that emphasising the affective properties of webpages decreases the number of negative webpages, and increases the number of positive webpages, participants expose themselves to. Clearly, we are not suggesting that one should make information consumption decisions based only on affective properties. To the contrary, we have written extensively about the multi-features of information critical in making information-consumption decisions, of which affect is only one (e.g., Sunstein & Sharot, 2020; Kelly & Sharot, 2021; Cogliati-Dezza et al., 2022; Charpentier et al., 2018; Vellani et al., 2020; Vellani et al., 2022). Instrumental utility of information and uncertainty reduction, for example, do and should drive information-seeking. What we envision is that affective labels could be used in the future together with other labels (such as the instrumental utility of information and its reliability) to empower users to make better information-consumption decisions that align with their goals.

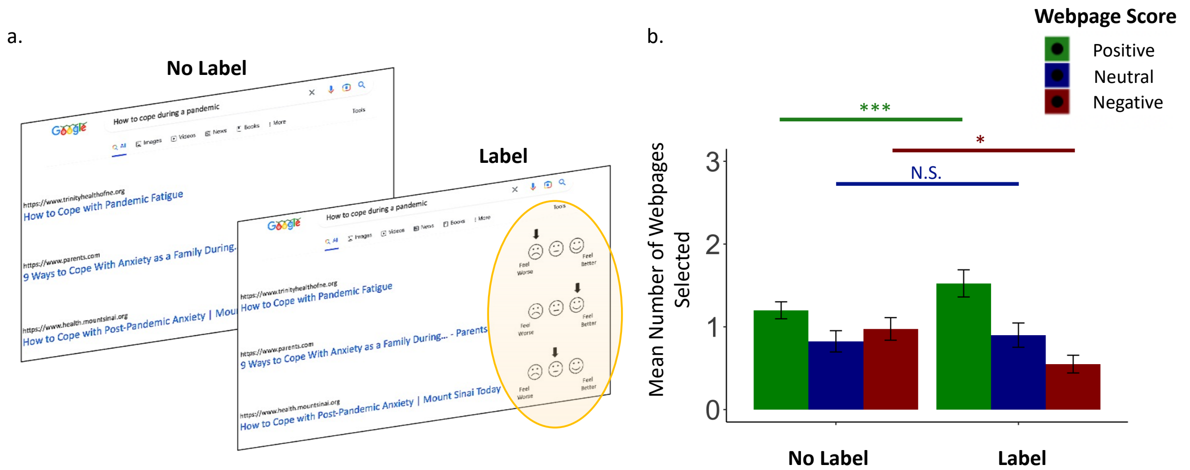

Figure 3.6. Novel online intervention decreases the amount of negative information browsed online. (a) Participants were assigned to either a label or no label condition. In the label condition they were presented with three Google search results pages from a set of 18. Each included three possible webpage links. Participants were asked to select the webpage they wanted to visit. In the label condition they also observed a label next to each link: either positive ('feel better'; green), neutral (blue), or negative ('feel worse'; red). The yellow oval is for illustrations purposes only and *was not* present in the actual study. **(b)** Participants in the label condition clicked on more webpages with the positive label and less webpages with the negative label than the no label condition. There was no difference in the number of neutral webpages selected. This suggests that cues indicating the effective properties of webpages alters participants web browsing patterns, such that they

expose themselves to less negative and more positive information. Error bars = standard error (SEM). *** = P < 0.001, * = P < 0.05, N.S. = not significant (two-sided).

3.4 Discussion

Our findings reveal that web-browsing both reflects and shapes mental health. The valence of the information people browse online was associated with their mental health, with those consuming more negative information tending to report worse mental health as measured by mood and self-reported psychopathology symptoms. A central question is whether browsing patterns alter mental health or vice versa. Our results support a reciprocal causal relationship between the affective properties of information consumed from self-guided searches and mood.

In particular, we show that participants who reported worse mood prior to browsing tended to access more negative content online. Exposure to negative content was in turn associated with worse mood post browsing (controlling for pre-browsing mood). We established the causality of this relationship by exposing participants to either negative or neutral webpages. We found that exposure to negative webpages resulted in worse mood, and this change in mood then led to more browsing of negatively valence information. Together, these findings reveal a feedback loop; low mood leads to the consumption of more negative information online, which in turn leads to worse mood and so on. However, it should be noted that, given the temporal dynamics of mood, an individual's mood can return to its previous state (i.e., pre-browsing mood; Zillmann, 2003) when starting a new browsing session. This could be a result of many factors, such as mood regulation strategies offline (Gross, 2002), or simply because we do not have the temporal resolution to assess how one might regulate their emotion online. For example, mood management theory (Zillmann, 1988) suggests that individuals are almost always exerting personal agency in shaping digital experiences and overall well-being; however, this agency may be diminished in people with poor well-being.

Our study is innovative in its approach of examining the link between the information browsed online and mental health. Previous research in this area has focused on analysing specific search engine queries rather than the actual text on webpages visited (Ayers et al., 2021; Gunnell et al., 2015). This traditional approach monitors certain keywords, such as "therapist" or "Prozac", to infer changes in population mental health. This method may be limited in its ability to assess an individual's mental health, as it only provides a limited dataset based on a few keywords, which would be used by individuals who are already aware of their symptoms and seek help.

The rational of our approach, namely to quantify the affective properties of text, is consistent with studies showing a relationship between the affective properties of shared content and mental health (i.e., such as posting on social media; De Choudhury et al., 2013; Chancellor et al., 2019; Kelley & Gillan, 2022; Eichstaedt et al., 2018). This may indicate an intriguing overlap between the mechanism governing information-seeking and those governing information-sharing (e.g., Vellani et al., 2022). Indeed, the size of the effects reported here are comparable to those found for the relationship between mental health and information-sharing (Kelley & Gillan, 2022) as well as for those found between mental health and frequency of social media use (Nan et al., 2024; Twenge et al., 2017; Yoon et al., 2019; McCrae et al., 2017; Vahedi & Zannella, 2019). An advantage of the current approach, is that it does not necessitate that people share information online, an activity that is clearly prevalent, but less so than online information-seeking.

The current results point to a consistent relationship between mental health and the consumption of less negatively valenced text, rather than more positively valanced text. This observation is consistent with past studies that find that individuals with poor mental health are more likely to attend to negative information (Fox, 1994; Roiser et al., 2012) and use it to alter their beliefs (Garrett & Sharot, 2014). It is possible, that the observed non-significant relationship between mental health and a tendency to consume positive information from self-guided searches reflects a ground truth, or alternatively, a small effect may exist that was not picked up by our data. For example, if participants intentionally adjusted their behaviour (i.e., not including all the webpages they actually visited or not sincerely completed the mental health questionnaires), this would induce noise that would make smaller effects difficult to detect. Steps can be taken in the future to enhance the methodology used to increase the likelihood of detecting such effects. Improvements can include adding analysis of images and videos; collecting timestamps of participants' web-browsing to measure the exact amount of time users spend on each piece of content; including password-protected websites such as social media platforms; including browsers beyond Firefox, and extending the duration of data collection to weeks or months while employing ecological momentary assessment (EMA). The latter will allow us to characterize the relationship between mental health and web-browsing patterns at a more granular level (e.g.., by applying time-series analysis). Finally, we used the NRC lexicon to assess valence because it showed the strongest relationship to human raters compared to other lexicons and compared to a standard large language model (LLM) approach that uses a transformers based architecture, even though the former does not incorporate context.

It is also interesting to consider the role of search algorithms here. Many algorithms are trained on participants' past behaviour, thus might perpetuate a participant's affective state by promoting specific types of information (such as negatively valenced content), potentially exacerbating the feedback loop identified. Moreover, algorithms may also be a source of noise, in the sense that they alter peoples' natural search intentions. If that is the case, the relationship between mental health and self-driven searches is likely even greater than we report here.

Given that Study 3 established a bi-directional relationship between exposure to negative information and mental health, we examined whether individuals would choose to access less negative information if they were made aware of the potential emotional impact of webpages before browsing. Indeed, our results showed that providing individuals with cues about the emotional impact of webpages effectively changed their browsing patterns, leading to a decrease in exposure to negative content and an increase in exposure to positive content, which in turn improved their mood. This result suggests that a simple intervention is effective in reducing exposure to negative information and potentially improving mood.

In many cases it would be obviously suboptimal to solely base information-consumption decisions on the affective properties of information. For example, if someone searches for information on whether smoking causes cancer, the most positive link may not necessarily be the wise choice. Thus, we do not envision the intervention described here as a stand-alone tool. Rather, by providing users with affective labels in addition to other labels, such as the reliability and instrumental utility of information, users can make more informed decisions that align with their current goals. For instance, users may want to prioritize the instrumental utility of information in one situation and prioritize their mood in another by focusing on affective labels. As such, our study not only provides evidence for the relationship between information-seeking and mental health, but can inform the development of tools aimed at enhancing mental health by improving information consumption decisions.

Our approach combines psychological theory with computer science, advancing theoretical understanding and the development of practical tools. It introduces a novel methodology — analysing web-content browsed — to explore the causal b-directional relationship between mental health and web-browsing patterns. The empirical findings also feed into the development of tools that can help users browse the web in an informed manner that can improve mood.

3.5 References

Ayers, J. W., Poliak, A., Johnson, D. C., Leas, E. C., Dredze, M., Caputi, T., & Nobles, A. L. (2021). Suicide-Related Internet Searches During the Early Stages of the COVID-19 Pandemic in the US. JAMA Network Open, 4(1). https://doi.org/10.1001/JAMANETWORKOPEN.2020.34261

Chancellor, S., Birnbaum, M., Caine, E., Silenzio, V. & De Choudhury, M. A taxonomy of ethical tensions in inferring mental health states from Social media. in Proc. Conference on Fairness, Accountability, and Transparency (FAT*) (ACM, 2019).

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, 115(31), E7255-E7264.

DataReportal (2022), “Digital 2022 Global Digital Overview,” retrieved from https://datareportal.com/reports/digital-2022-global-overview-report

De Choudhury, M., Gamon, M., Counts, S. & Horvitz, E. (2013). Predicting depression via social media. ICWSM 2, 128–137.

Devlin, J., Chang, M. W., Lee, K., & Toutanova, K. (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. arXiv preprint arXiv:1810.04805.

Dezza, I. C., Maher, C., & Sharot, T. (2022). People adaptively use information to improve their internal states and external outcomes. Cognition, 228, 105224.

Eichstaedt, J. C. et al. Facebook language predicts depression in medical records. Proc. Natl Acad. Sci. USA 115, 11203–11208 (2018).

Ertel, K. A., Glymour, M. M., & Berkman, L. F. (2009). Social networks and health: A life course perspective integrating observational and experimental evidence. Journal of Social and Personal Relationships, 26(1), 73-92.

Ettman, C. K., Adam, G. P., Clark, M. A., Wilson, I. B., Vivier, P. M., & Galea, S. (2022). Wealth and depression: A scoping review. Brain and Behavior, 12(3), e2486.

Gillan, C. M., Kosinski, M., Whelan, R., Phelps, E. A., & Daw, N. D. (2016). Characterizing a psychiatric symptom dimension related to deficits in goal-directed control. elife, 5, e11305.

Goldberg, Y., & Levy, O. (2014). word2vec Explained: deriving Mikolov et al.'s negative-sampling word-embedding method. arXiv preprint arXiv:1402.3722.

Gunnell, D., Derges, J., Chang, S. Sen, & Biddle, L. (2015). Searching for suicide methods. Crisis, 36(5), 325–331. https://doi.org/10.1027/0227-5910/A000326

Hu, M., & Liu, B. (2004, August). Mining and summarizing customer reviews. In Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining (pp. 168-177).

Kelley, S. W., & Gillan, C. M. (2022). Using language in social media posts to study the network dynamics of depression longitudinally. Nature communications, 13(1), 870.

Kelly, C. A., & Sharot, T. (2021). Individual differences in information-seeking. Nature communications, 12(1), 7062.

Kiritchenko, S., Hipson, W. E., Coplan, R. J., & Mohammad, S. M. (2020). SOLO: A corpus of tweets for examining the state of being alone. arXiv preprint arXiv:2006.03096.

Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M. & Gottlieb, J. Diverse motives for human curiosity. Nature Human Behaviour. 3, 587–595. https://doi.org/10.1038/s41562-019-0589-3 (2019).

Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., ... & Stoyanov, V. (2019). Roberta: A robustly optimized bert pretraining approach. arXiv preprint arXiv:1907.11692.

Marconcin, P., Werneck, A. O., Peralta, M., Ihle, A., Gouveia, É. R., Ferrari, G., ... & Marques, A. (2022). The association between physical activity and mental health during the first year of the COVID-19 pandemic: a systematic review. BMC Public Health, 22(1), 209.

Mohammad, S. (2018). Obtaining reliable human ratings of valence, arousal, and dominance for 20,000 English words. In Proceedings of the 56th annual meeting of the association for computational linguistics (volume 1: Long papers) (pp. 174-184).

Mohammad, S. M., & Turney, P. D. (2013). Crowdsourcing a word–emotion association lexicon. Computational intelligence, 29(3), 436-465.

Peluso, M. A. M., & De Andrade, L. H. S. G. (2005). Physical activity and mental health: the association between exercise and mood. Clinics, 60(1), 61-70.

Pieh, C., O´ Rourke, T., Budimir, S., & Probst, T. (2020). Relationship quality and mental health during COVID-19 lockdown. Plos one, 15(9), e0238906.

Pollack, C. E., Chideya, S., Cubbin, C., Williams, B., Dekker, M., & Braveman, P. (2007). Should health studies measure wealth?: A systematic review. American journal of preventive medicine, 33(3), 250-264.

Robles, T. F., & Kiecolt-Glaser, J. K. (2003). The physiology of marriage: Pathways to health. Physiology & behavior, 79(3), 409-416.

Rouault, M., Seow, T., Gillan, C. M., & Fleming, S. M. (2018). Psychiatric symptom dimensions are associated with dissociable shifts in metacognition but not task performance. Biological psychiatry, 84(6), 443-451.

Seow, T. X., & Gillan, C. M. (2020). Transdiagnostic phenotyping reveals a host of metacognitive deficits implicated in compulsivity. Scientific reports, 10(1), 1-11.

Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. Nature Human Behaviour, 4(1), 14-19.

Sunstein, C. R. (2019). Ruining popcorn? The welfare effects of information. Journal of Risk and Uncertainty, 58, 121-142.

VandenBos, G. R. (2007). APA dictionary of psychology. American Psychological Association.

Vellani, V., & Sharot, T. (2022). How people decide when to inform others.

Vellani, V., de Vries, L. P., Gaule, A., & Sharot, T. (2020). A selective effect of dopamine on information-seeking. elife, 9, e59152.

Yazman, J. (2017). A pair of text analysis explorations. Retrieved August 17, 2022, from https://joshyazman.github.io/pair-of-text-analysis-explorations/

Lerner, J. S., Li, Y., Valdesolo, P., & Kassam, K. S. (2015). Emotion and decision making. Annual review of psychology, 66, 799-823.

Hockey, G. R. J., John Maule, A., Clough, P. J., & Bdzola, L. (2000). Effects of negative mood states on risk in everyday decision making. Cognition & Emotion, 14(6), 823-855.
George, J. M., & Dane, E. (2016). Affect, emotion, and decision making. Organizational Behavior and Human Decision Processes, 136, 47-55.
Paulus, M. P., & Angela, J. Y. (2012). Emotion and decision-making: affect-driven belief systems in anxiety and depression. Trends in cognitive sciences, 16(9), 476-483.
Phelps, E. A., Lempert, K. M., & Sokol-Hessner, P. (2014). Emotion and decision making: multiple modulatory neural circuits. Annual review of neuroscience, 37, 263-287.
Pictet, A., Coughtrey, A. E., Mathews, A., & Holmes, E. A. (2011). Fishing for happiness: The effects of generating positive imagery on mood and behaviour. Behaviour research and therapy, 49(12), 885-891.
MacLeod, C., & Mathews, A. (1988). Anxiety and the allocation of attention to threat. The Quarterly journal of experimental psychology, 40(4), 653-670.
Lyubomirsky, S., Caldwell, N. D., & Nolen-Hoeksema, S. (1998). Effects of ruminative and distracting responses to depressed mood on retrieval of autobiographical memories. Journal of personality and social psychology, 75(1), 166.

Chapter 4

High-Level Characteristics of Web Searches Change Under Stress

4.1 Overview

Every person will face unexpected adversities during their lifetime. These stressful events may be global (e.g., war, pandemic) or unique to the individual (e.g., being diagnosed with cancer, losing one's job, divorce). Abundant research highlights that such events often lead to stress, anxiety, confusion, and a reduced sense of control, impacting mental and emotional well-being (Finlay-Jones & Brown, 1981; Francis et al., 2012; McLaughlin et al., 2010; Miloyan et al., 2018; Suls & Mullen, 1981; Globig et al., 2020). The American Psychological Association (APA; VandenBos, 2007) defines a stressful event as "*an occurrence or circumstance that individuals perceive as threatening, challenging, or demanding, thereby eliciting a stress response.*" This broad definition encapsulates a range of experiences, from major life changes and daily hassles to situations such as job loss, financial difficulties, relationship conflicts, health issues, traumatic experiences, and environmental disasters (VandenBos, 2007).

One available adaptive reaction to stress is to seek information that can help guide action to promote adaptation (Hirshleifer & Riley, 1979; Sharot & Sunstein, 2020; Stigler, 1961).Such actions can be directly related to the event experienced (e.g., during wartime people may search for information on how to secure windows from being shuttered by rockets) or indirectly related (e.g., searching for activities that can distract oneself from the adversity). However, it's important to note that this theory of information-seeking as an adaptive strategy is contingent on the individual's sense of agency in the situation. In scenarios where the individual perceives no control or ability to influence the outcome, as in the case of an unavoidable and immediate stressor like being suddenly immersed in an ice bath, seeking information may not be considered a viable or useful response.

To date, research on the relationship between information-seeking and stress has mostly focused on the frequency of information-seeking. Some studies propose that stress is associated with greater information seeking (Drouin et al., 2020; Ebrahim et al., 2020; Loosen et al., 2021), which may decrease the sense of uncertainty that is heightened under stress. Others, however, suggest that stress leads to an avoidance reaction which is characterised by less information-seeking about the stressor (Kash et al. 2000; Chae, 2016).

Here, we take a different theoretical viewpoint. Rather than focusing on whether stress generally enhances or reduces information seeking, we test the hypothesis that when experiencing abrupt stressful life events people are more likely to search for information that can direct action - a reaction which may be adaptive. In other words, stress may alter the type, rather than frequency, of information people seek. Across multiple studies, we test whether such changes can be detected and quantified by simple analysis of web searches. We propose that by examining the features of the information that people seek, we can gain insight into their external state and their internal reaction to that state.

To that end, we also examine whether the negative context people find themselves in is associated with a change in the valence of web searches. That is, whether a negative state may lead to more negatively valanced searches due to that state, or alternatively to more positive searches in an attempt perhaps to counter the negative state. The former possibility is supported by studies showing that anxious individuals have a bias towards negative stimuli (Bar-Haim et al., 2007; Cisler et al., 2010; MacLeod et al., 1988), which suggests that anxiety may increase the search for negative information.

We test our proposal both in the context of a global stressful event (i.e., the COVID-19 pandemic) and of personal aversive life events. Our approach differs dramatically from past attempts to relate web searches to mental state. Current research on web searches has focused on relating web searches for specific content terms (e.g., '*suicide'*, '*anxiety*', *'Prozac'*) with mental health indicators of a population, a method that has resulted in mixed findings (Ayers et al., 2021; Gunnell et al., 2015; Sueki, 2011; Hoerger et al., 2020; Barros et al., 2019; Ayers et al., 2012; Knipe et al., 2020; McCarthy, 2010; Misiak et al., 2020; Rana, 2020; Sinyor et al., 2020; Tran et al., 2017; Arora et al., 2019; Yang et al., 2010, Suh et al., 2021). In contrast, we assess whether changes in the *high-level features* of searches reflects stress levels. This approach, which is based on a recent theory of information-seeking motives (Sharot & Sunstein, 2020), may be more sensitive as it does not make an assumption about which topics people are searching for, but rather the characteristic of information they are searching for (i.e., information that may guide action).

First, we conducted a control study to identify which question-words people would use when seeking information to guide action (Study 1). The results clearly show that participants selectively use *"How"* for this purpose. Next in Study 2, we calculated the percentage of Google searches containing the *"How"* question-word submitted in the UK and US out of all searches submitted in that region, every week for a year from the date a *"National Emergency"* was declared and compared to the years before. Changes to the proportion of *"How"* searches cannot be explained by changes in the volume of Google searches during the pandemic, as we examined the change in the *percentage* of *"How"* searches out of all searches at that time. We also quantified the valence of the most frequent questions submitted to the Google search engine each week in the UK and US using a machine learning approach (HuggingFace, 2022). We then examined how these features related to weekly stress reports of approximately 17K individuals in the UK. Importantly, we dissociate the effects of stress on information-seeking from the effect of COVID-19 related confinement. Together, these analyses enable us to examine how features of web searches alter under a global stressor.

To be able to generalise our findings to private stressful events and to rule out potential third factors, we then tested our hypotheses in a controlled environment in Study 3. Here, we manipulate stress levels and examined if the manipulated stress impacts the likelihood of asking *"How"* questions in relation to private events. If successful, our approach of quantifying high-level features of *web searches* and relating them to stress may provide a new avenue for monitoring population-level stress during times of crisis.

4.2 Methods

Study 1

Participants. One-hundred participants (Age =39.29 (SD = 13.86), Females = 56%, Males = 44%, Other = 0%) completed the study on Qualtrics (www.qualtrics.com) and were recruited via Prolific's online recruitment system (www.prolific.co). Participants received £7.50 per hour for their participation. For all studies presented in this article, ethical approval has been provided by the Research Ethics Committee at UCL and all participants have given their informed consent to participate. All methods were performed in accordance with UCL's guidelines and regulations.

Procedure. Participants were asked to think about a goal they were trying to achieve. They were then instructed to select from a list of eight question-words (*"What"*, *"Which"*, *"Who"*, *"Where"*, *"Why"*, *"When"*, *"Whose"*, and *"How"*) a word that they would use to submit a query on Google to help guide their actions to achieve their goal (i.e., experimental condition). In the control condition participants were asked to think about a topic they are interested in learning more about and select a question-word from the same list to ask a query on Google to increase their knowledge about that topic. The order of the two questions were counter balanced, and the order of the question-words were randomised. This design allowed us to identify which words people use to seek information to guide action and test whether such words were used generally for information-seeking and more specifically to guide action.

Analysis. To assess which question-words were associated with guiding action, we first calculated the proportion of people that selected each of the question-words. We then conducted a Fisher's exact test to examine whether the most prevalent question-word selected (i.e., *"How")* was significantly different than the proportion of all other question-words asked together for each condition separately (coded '0' if *"How"* was selected and '1' if any other question-word was selected), and whether the proportion of people that selected *"How"* in the control condition was significantly different than in the experimental condition.

Study 2

Data Extraction

Web Search Data. Weekly search data was extracted from Google Trends (www.googletrends.com) for 220 weeks (January 1st, 2017, through March 21st, 2021). This was done separately for the UK and the US. Based on the results of Study 1, to quantify action guidance, we extracted the Google search volume index for the search term *"How".* We also extracted the Google search volume index for the search terms *"What"*, *"Which"*, *"Who"*, *"Where"*, *"Why"*, *"When"*, and *"Whose"* and then averaged them together to quantify general question asking (i.e., control variable). A Google search volume index value is equal to the number of searches for the specific term of interest in a given week and region (for example total number of searches that include the question-word *"How"* in the UK the first week of 2020) divided by the total number of searches in that same time and region (for example the total number of Google searches submitted in UK the first week of 2020). These values are normalised to represent search interest relative to the highest value for that region for the entire time frame (i.e., January 1st, 2017 – March 21st, 2021).

To quantify valence, we extracted the 25 most popular search searches for each week and region for *"What"*, *"Which"*, *"Who"*, *"Where"*, *"Why"*, *"When"*, *"Whose"* and "*How",* questions. That is, for each week we extract up to 25 search searches per type of question (i.e., ~200 total web searches for each week), as this is the maximum Google Trends reports. Next, we implemented a machine learning approach to assess the valence of these searches, by applying the pre-trained model, distilbert-base-uncased-finetuned-sst-2-english (HuggingFace, 2022), to the extracted data. The models output contains two labels: positive and negative, along with a score between 1 and 0 (1 = absolute confidence in the model output label, 0 = zero confidence in the model output label). The sum of the values is always equal to 1, for example, a potential output could be Positive = 0.9, Negative = 0.1. As the two measures are fully dependent, we used the positive confidence score as our measurement of valence and then transformed this number to be on a scale from 0 to 100, such that it would be easily comparable to the Google search volume index for *"How"*, described above (i.e., a score of 100 denotes the most positively valenced score and a score of 0 the most negatively valenced score).

Self-Reported Stress. Data was extracted, with permission, from the UK COVID-19 Social Study (Fancourt et al., 2021). The study is a panel study of over 70,000 UK citizens which aims to characterise the psychological and social experience of adults living in the UK during the Covid-19 pandemic (**see Table 4.1 for demographics**). The study commenced as a weekly survey, with participants receiving an invitation to the next wave of data collection 7 days following their last completion. All participants received up to 2 reminders (24 and 48 hours following their initial weekly invitation). The link to their last reminder remained live so they could return to the study a few days later if they chose to. Following week 22 of the study, monthly follow-ups rather than weekly follow-ups were sent. To attain an equal number of responses across time, participants were randomised to receive their monthly invitation on either week 1,2,3 or 4 of the month, with subsequent invitations following 28 days after they completed the survey. An average of 17,468 individuals submitted data each week (**see Table 4.2 for response frequency for each week**). For full methods and demographics for the sample see www.COVIDSocialStudy.org. The UK COVID-19 Social Study was approved by the UCL Research Ethics Committee, and all participants gave written informed consent. All methods were performed in accordance with UCL's guidelines and regulations.

Participants were asked: "*over the past week, have any of the following been worrying you at all, even if only in a minor way?"* They were presented with 18 factors that may cause worry (for example internet access, boredom, neighbours) and were to pick any that they were worried about. Five of these factors were a-priori categorised by the authors of the survey (Fancourt et al., 2021) as ones that have been impacted by COVID. These were (i) catching Covid-19 (ii) becoming seriously ill from Covid-19, (iii) finances, (iv) losing your job/unemployment and (v) getting food. Second, they were asked *"have any of these things been causing you significant stress? (e.g., they have been constantly on your mind or have been keeping you awake at night)"*. They were presented with the same 18 factors as above and were asked to tick any of those causing significant stress. For each week and factor, Fancourt et al., 2021 calculated the proportion of respondents that ticked that factor either in response to question 1 and/or question 2. Factors i and ii were a-priori combined by Fancourt and colleagues (2021) to make one factor, leaving us with four factors. For each week the proportion of people ticking 1 and/or 2 were averaged across the four factors to produce one indicator of 'stress levels' for that week.

Table 4.1 shows the demographic of respondents to the UK COVID-19 Social Study. Importantly, data points reported by Fancourt et al., (2021) were weighted using auxiliary weights to the national census and Office for National Statistics (ONS) data. We used these weighted data points in our study. Thus, reported stress levels are representative of the UK population.

Table 4.1. *Demographics of respondents in the UK COVID-19 Social Study (adapted from Fancourt et al., 2021). Data in the Fancourt et al., (2021) study and in our study* ***are weighted using auxiliary weights*** *to the national census and Office for National Statistics (ONS) data.*

		Number of observations	%			Number of observations	%
Age				**Education levels**			
	18-29	51,858	5.77		GCSE or below	126,427	14.1
	30-59	493,016	54.9		A-levels of equivalent	154,954	17.3
	60+	353,559	39.4		Degree or above	617,052	68.7
Gender				**Any diagnosed mental health conditions**			
	Male	225,578	25.2		No	748,416	83.3
	Female	669,279	74.8		Yes	150,017	16.7
Ethnicity				**Any diagnosed physical health conditions**			
	White	860,157	96.0		No	516,884	57.5
	Ethnic minority	35,455	3.96		Yes	381,549	42.5
UK nations				**Keyworker**			
	England	725,705	81.6		No	711,201	79.2
	Wales	108,598	12.2		Yes	187,232	20.8
	Scotland	55,416	6.23	**Living with children**			
Living arrangement					No (excluding those who live alone)	510,650	72.0
	Not living alone	709,289	79.0		Yes	198,639	28.0
	Living alone	189,144	21.1	**Living area**			
Annual household income					Village/hamlet/isolated dwelling	225,022	25.1
	>30k	482,268	59.6		City/large town/small town	673,411	75.0
	<30k	327,187	40.4				

Table 4.2. *The total number of participants providing data during each calendar week in the UK COVID-19 Social Study (adapted from Fancourt et al., 2021).*

Date	Week	Freq	Date	Week	Freq
21/03/20-27/03/20	1	28,929	**19/09/20-25/09/20**	27	8,318
28/03/20-03/04/20	2	27,873	**26/09/20-02/10/20**	28	8,366
04/04/20-10/04/20	3	38,151	**03/10/20-09/10/20**	29	8,501
11/04/20-17/04/20	4	38,453	**10/10/20-16/10/20**	30	8,072
18/04/20-24/04/20	5	38,504	**17/10/20-23/10/20**	31	7,495
25/04/20-01/05/20	6	36,513	**24/10/20-30/10/20**	32	7,612
02/05/20-08/05/20	7	36,651	**31/10/20-06/11/20**	33	7,830
09/05/20-15/05/20	8	37,549	**07/11/20-13/11/20**	34	7,443
16/05/20-22/05/20	9	35,702	**14/11/20-20/11/20**	35	6,995
23/05/20-29/05/20	10	33,293	**21/11/20-27/11/20**	36	7,078
30/05/20-05/06/20	11	32,196	**28/11/20-04/12/20**	37	7,190
06/06/20-12/06/20	12	31,304	**05/12/20-11/12/20**	38	6,947
13/06/20-19/06/20	13	30,229	**12/12/20-18/12/20**	39	6,473
20/06/20-26/06/20	14	29,153	**19/12/20-25/12/20**	40	6,240
27/06/20-03/07/20	15	28,534	**26/12/20-01/01/21**	41	6,966
04/07/20-10/07/20	16	27,552	**02/01/21-08/01/21**	42	7,038
11/07/20-17/07/20	17	26,737	**09/01/21-15/01/21**	43	6,274
18/07/20-24/07/20	18	25,983	**16/01/21-15/01/21**	44	6,219
25/07/20-31/07/20	19	25,005	**23/01/21-29/01/21**	45	6,540
01/08/20-07/08/20	20	24,530	**30/01/21-05/02/21**	46	6,831
08/08/20-14/08/20	21	23,851	**06/02/21-12/02/21**	47	6,048
15/08/20-21/08/20	22	23,120	**13/02/21-19/02/21**	48	6,217
22/08/20-28/08/20	23	11,373	**20/02/21-26/02/21**	49	6,111
29/08/20-04/09/20	24	10,025	**27/02/21-05/03/21**	50	6,574
05/09/20-11/09/20	25	9,916	**06/03/21-12/03/21**	51	8,683
12/09/20-18/09/20	26	10,009	**13/03/21-19/03/21**	52	9,128

COVID-19 Confinement Score. To measure COVID-19 related confinement, we extracted eight confinement variables from a publicly available dataset (The Oxford University COVID-19 Government Response Tracker; Webster et al., 2021). All variables are ordinal coded by severity/intensity of confinement, on a daily basis (from January 1st, 2020 to March 21st, 2021), for the following: (i) *school and university closures*, (ii) *workplace closures*, (iii) *public event cancelations*, (iv) *restrictions on gatherings*, (v) *public transport restrictions*, (vi) *stay at home requirements*, (vii) *restrictions on domestic travel,* and (viii) *restrictions on international travel*; **see Table 4.3 for coding.** To obtain weekly values, we computed weekly averages of the daily ratings. To quantify an overall COVID-19 related confinement score, we transformed all variables to range between 0 and 1 using the R function *scaler* from the R package, *bruceR*. Finally, we averaged the 8 transformed variables together.

Table 4.3. *Coding of COIVID-19 related confinement variables (adapted from Webster et al, 2021).*

Variable Description	Coding instructions
Closings of schools and universities	0 - No measures 1 – Recommend closing, or all schools open with alterations resulting in significant differences compared to usual, non-Covid-19 operations 2 - Require closing (only some levels or categories, e.g., just high school, or just public schools) 3 - Require closing all levels
Closings of workplaces	0 - No measures 1 - recommend closing (or work from home) 2 - require closing (or work from home) for some sectors or categories of workers 3 - require closing (or work from home) all-but-essential workplaces (E.g., grocery stores, doctors)
Cancelling public events	0 – No measures 1 - Recommend cancelling 2 - Require cancelling
Cut-off size for bans on gatherings	0 - No restrictions 1 - Restrictions on very large gatherings (the limit is above 1000 people) 2 - Restrictions on gatherings between 101-1000 people 3 - Restrictions on gatherings between 11-100 people 4 - Restrictions on gatherings of 10 people or less
Closing of public transport	0 - No measures 1 - Recommend closing (or significantly reduce volume/route/means of transport available) 2 - Require closing (or prohibit most citizens from using it)
Orders to "shelter-in- place" and otherwise confine to home	0 - No measures 1 - recommend not leaving house 2 - require not leaving house with exceptions for daily exercise, grocery shopping, and 'essential' trips 3 - Require not leaving house with minimal exceptions (E.g., allowed to leave only once a week, or only one person can leave at a time, etc.)
Restrictions on internal movement	0 - No measures 1 - Recommend not to travel between regions/cities 2 – internal movement restrictions in place
Restrictions on international travel	0 - No measures 1 - Screening 2 - Quarantine arrivals from high-risk regions 3 - Ban on arrivals from some regions 4 – Ban on all regions or total border closure

Analysis

Analysis was conducted separately for data from the UK and the US. In each region we quantified action guidance searches (i.e., *"How"*) and valence of searches every week from the date the *"National Emergency"* was declared due to the COVID-19 pandemic, (March 23rd, 2020, in the UK and March 13th, 2020 in the US) through March 21st, 2021, as well as every week dating back to January 1st, 2017. We then compared the weekly scores before the *"National Emergency"* to that after using an independent samples t-test.

To assess whether our measures were related to UK stress levels, we conducted two linear models predicting on a weekly basis Google's search volume index of *"How"* questions and the Valence Index of questions submitted to Google in the UK, from UK stress levels. We also included weekly COVID-related confinement scores in the models to disentangle the effects of stress from the effects of confinement due to restrictions placed by the Government. To account for simple temporal trends, we removed the linear trend from the dependent and predictor variables first, using the *detrend* function in the *pracm*a R package. The detrended dependent and predictor variables were then Z-scored.

Next, to examine whether the relationship between stress and *"How"* searches was specific or rather reflected a general tendency to ask more questions when stressed we conducted a third linear model relating question asking to stress levels controlling for COVID-related confinement scores. To do so, we extracted the Google search volume index for all other common question-words (i.e., *"What", "Which", "Who", "Where", "Why", "When",* and *"Whose"*; Z-scored and detrended) and then averaged them together.

We were then interested in whether stress was better predicted by *"How"* Google searches than searches for stress related terms. To test this, we first removed the linear trend from the dependent and predictor variables, and then Z-scored the dependent and predictor variables. We then ran a model predicting the proportion of UK sample reporting COVID-related stress from the *"How"* Google search index as well as Google search index for the word's '*stress*', '*anxiety*', and '*mental health*' in the UK. In addition, we ran multiple linear models to predict UK COVID-related stress levels each time from only one of the terms above.

Finally, we tested the predictive validity of a simple model using stress levels to predict the proportion of *"How"* searches using a leave one out analysis. Once again, we removed the linear trend from the dependent and predictor variables first, and then the dependent and predictor variables were Z-scored. Specifically, the simple model was run on all the data save for one time point which was held out from the analysis. We then used the regression beta to predict the proportion of *"How"* searches of the left-out time point. This process was repeated so that each week's proportion of *"How"* searches was estimated from the simple model parameters generated without using that week to fit the data. This resulted in two values for the proportion of *"How"* searches for each week: the actual proportion of *"How"* searches (data) and the predicted value from the leave-one-out validation (estimate). The actual proportion of *"How"* searches of a week (data) and the predicted proportion of *"How"* searches (estimation) were then correlated and compared using a paired sample t-test. This analysis indicates whether the population stress levels is a good predictor of the proportion of *"How"* searches.

Study 3

Participants. One hundred and ninety-three participants completed the study on Qualtrics (*www.qualtrics.com*) and were recruited via Prolific's online recruitment (*www.prolific.co*). Participants received £7.50 per hour for their participation. One participant was excluded for not providing a valid response, leaving the final participant N at 192 (stress condition: n = 99, age = 39.45, SD = 14.74; Females = 78.8%, Males = 19.2%, Other – 2.0%; control condition: n = 93, age = 38.46, SD = 11.91; Females = 76.3%, Males = 22.6%, Other = 1.1%).

Procedure. Participants were asked to recall a time when they were very stressed (stress condition) or recall a time when they were happy and relaxed (*control condition*). They were then instructed to think about that time in as much detail as possible and describe it in a text box. they indicate their stress level on a scale ranging from very calm (-50) to very stressed (+50) before and after the induction. Next, participants were asked to enter two searches they could have entered to Google during the time.

Analysis. To assess whether the manipulation was successful, a 2 (condition: stress, control) by 2 (time: pre-induction, post induction) ANOVA was run with follow up pair-wise t-tests. Next, for each participant, we counted the number of questions that began with *"How"* (i.e., 0,1,2) and for every other question word (i.e., "*What", "Which", "Who", "Where", "Why", "When",* and *"Whose"*). We then conducted separate independent samples t-tests to assess the difference in the number asked for each question type between conditions. Finally, we implemented a machine learning approach to assess the average valence of participants' two searches by applying the pre-trained91 model distilbert-base-uncased-finetuned-sst-2-english (HuggingFace, 2022) to their submitted searches. Finally, we examined the difference between participants in the stress and control conditions with regard to the number of *"How"* searches submitted and the valence of their searches.

4.3 Results

Study 1

***"How"* web searches are selectively associated with the need to guide action.** To determine which question-words are associated with guiding action, we asked 100 participants (Age =39.29 (SD = 13.86), Females = 56%, Males = 44%, Other = 0%) to think about a goal they were trying to achieve. They were then instructed to select from a list of eight question words ("*What", "Which", "Who", "Where", "Why", "When", "Whose",* and *"How"*), the word that they would use in a Google query to help guide their actions to achieve the goal (i.e., experimental condition). In the control condition participants were asked to think about a topic they are interested in learning more about. They were then instructed to select a question-word from the same list of eight words, the one they would use in Google query to increase their knowledge about the topic. This design allowed us to identify which words people use to seek information to guide action and test whether such words were used generally for information-seeking, or more specifically to guide action.

The likelihood of selecting the word *"How"* to ask a question in the experimental condition (to guide action) was equal to 87%, which was significantly greater than the likelihood of selecting *"How"* in the control condition (to increase understanding), which was 29%. ($p = 0.009$, Fisher exact test **see Figure 4.1**). Moreover, in the experimental condition, the likelihood of selecting "*How*" was significantly greater than the likelihood of selecting all other question words put together ($p < 0.001$, Fisher exact test; **see Figure 4.1**), while in the control condition, the likelihood of selecting "*How*" was significantly less than the likelihood of selecting all other question words put together ($p< 0.001$, Fisher exact test; **see Figure 4.1**).

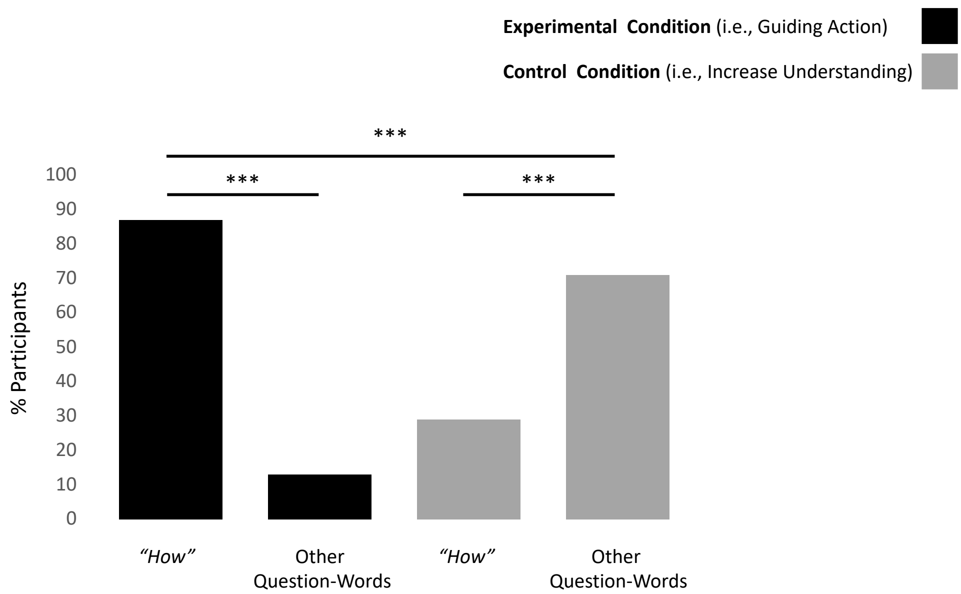

Figure 4.1. ***"How"* questions are associated with guiding actions.** Plotted on the y-axis is the percentage of participants selecting a particular question-word. Participants were more likely to select *"How"* over other question-words when asking a question to help guide their actions to achieve a goal. They also were more likely to ask *"How"* to help guide their actions than to simply increase their understanding. *** = $P < 0.001$, ** = $P < 0.01$.

Given the results of Study 1, we used "*How*" questions in Study 2 and 3 as a proxy for the desire to gain knowledge that could help guide action. Note that we are not suggesting that all questions that begin with "*How*" are intended to guide action. Rather that, on average, if people want to ask a question to guide action, they will be likely to use "*How*". In contrast, they are not especially likely to use "*How*" simply to learn more about a topic.

Study 2

The pandemic resulted in a significant change to the high-level features of web searches. To assess the high-level feature of web searches associated with guiding action, motivated from Study 1, we extracted the Google search volume index of "*How*" questions. A Google search volume index is equal to the number of searches for the specific term of interest in a given week and region divided by the total number of searches for that same week and region. These percentages are then normalised to represent search interest relative to the highest percent for that region for the entire time frame (i.e., January 1st, 2017 - March 21st, 2021; **see Method for details**). Note, that weekly changes to the Google search volume index cannot be explained by weekly changes in the total volume of Google searches, as the index reflects the *percent* of specific searches out of all searches that week. We also calculated a second feature - a Valence Index - which indicates the valence of the most frequent questions submitted to Google search engine. To calculate this index, we first extracted the most frequent web searches each week that included question-words (i.e., *"What" "Which"*, *"Who"*, *"Where"*, *"Why"*, *"When"*, *"Whose"*, and *"How")*. Next, we implemented a machine learning approach to assess the valence of searches, by applying the pre-trained model distilbert-base-uncased-finetuned-sst-2-english (HuggingFace, 2022) to the extracted data. The models output contains a label: positive and negative, along with a score between 1 and 0 (1 = absolute confidence in the model output label, 0 = zero confidence in the model output label). As the sum of both values is equal to 1, we used only the positive confidence score as our measurement of valence and then transformed this number to be on a scale from 0 to 100, such that it would be easily comparable to the Google search volume index for *"How"*, described above (i.e., a score of 100 denotes the most positively valenced score and a score of 0 the most negatively valenced score). We validated the algorithm score in this context by asking a naïve human subject to categorise 192 randomly sampled searches as either 'more positive' or 'more negative'. We found a significant positive association between the algorithmic score and the human score ($r(190) = 0.463$, $p < 0.001$).

Analyses were conducted separately in the UK and the US. In each country we quantified the measures above for every week from the date the *"National Emergency"* was declared (UK: March 23rd, 2020; and US: March 13th, 2020) through March 21st, 2021, as well as every week from January 1st, 2017, until *"National Emergency"* was declared. There was a significant increase in the Google search volume index of *"How"* questions following the declaration of a *"National Emergency"* relative to the three years previous (*"How"* UK: before *"National Emergency"* declared: $M = 65.37$, $SD = 3.27$, after *"National Emergency"* declared: $M = 84.19$, $SD = 8.52$, $t(55.71) = -15.585$, $p < 0.001$, Cohen's $d = -3.754$; *"How"* US: before *"National Emergency"* declared: $M = 71.65$, $SD = 2.87$, after *"National Emergency"* declared: $M = 85.62$, $SD = 4.90$, $t(63.67) = -19.682$, $p < 0.001$, Cohen's $d = -4.029$; **Figure 4.2a-d**).

Moreover, the Valence Index after the *"National Emergency"* was declared was found to be more negative to before (Valence Index UK: before *"National Emergency"* declared: $M = 59.37$, $SD = 16.71$, after the *"National Emergency"* declared: $M = 49.77$, $SD = 13.26$, $t(218) = 3.788$, $p < 0.001$, Cohen's $d = 0.601$; Valence Index US: before

“National Emergency” declared: M = 64.09, SD = 14.52, after *“National Emergency”* declared: M = 57.07, SD = 17.41, t(218) = 2.917, p = 0.004, Cohen’s d = 0.460; **Figure 4.2e-h**).

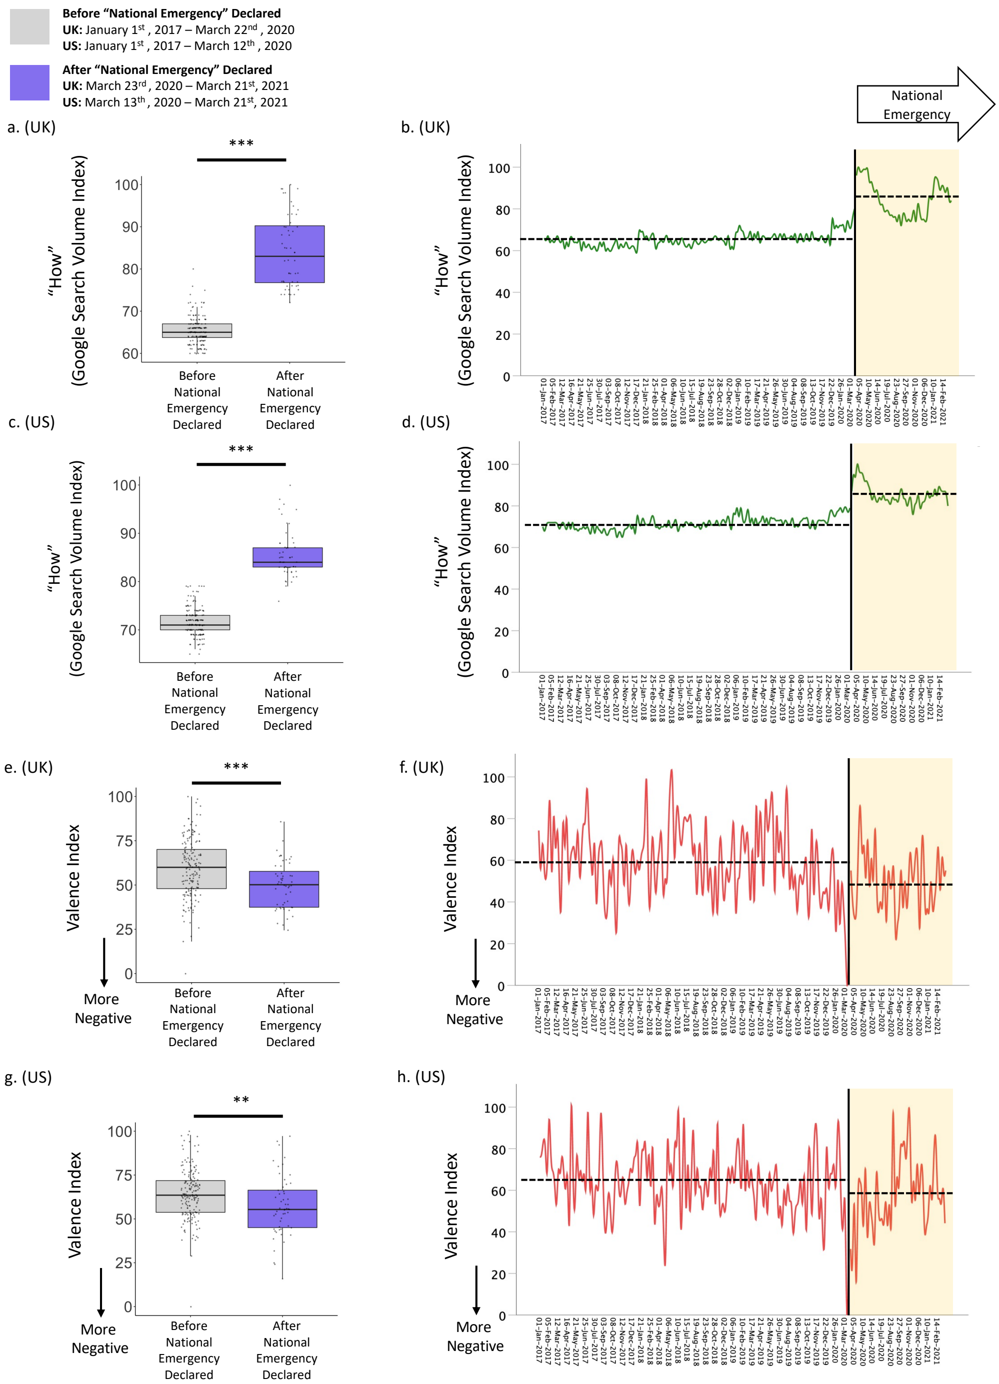

Figure 4.2. High-level characteristics of web searches alter during the pandemic. Relative volume of *“How”* searches (i.e., the proportion of “*How*” searches relative to all searches for that time and place) was greater after the COVID-19 *“National Emergency”* declaration relative to before in **(a&b)** the UK and **(c&d)** the US. The Valence Index [0 (most negative valenced) to 100 (most positive valenced)] reveals that searches submitted to the Google search engine were more negatively valenced

in the **(e&f)** UK and **(g&h)** the US after the COVID-19 *"National Emergency"* declaration relative to before. The period assessed prior to the *"National Emergency"* was from January 1st, 2017, to the declaration of each country's *"National Emergency"*. The *"National Emergency"* was assessed from March 23rd, 2020 to March 21st, 2021, in the **(a,b,e&f)** UK and from Match 13th, 2020 to March 21st, 2021 in the **(c,d,g&h)** US. **(a,c,e&g)** The horizontal lines indicate median values, boxes indicate 25–75% interquartile range and whiskers indicate 1.5 × interquartile range; individual scores are shown as dots. **(b,d,f&h)** The bold line indicates the declaration of the *"National Emergency"*, the dashed lines indicate the mean values for before and after the *"National Emergency"*. *** = $P < 0.001$, ** = $P < 0.01$ (two-sided).

Population stress-levels are selectively associated with asking *"How"*. Thus far, we have shown that there is an increase in the proportion of *"How"* and negatively valenced searches submitted to the Google search engine during the pandemic relative to before. We next examined whether these were related to population stress levels. We had access to self-report stress levels collected every week in the UK between March 21st, 2020, and March 21st, 2021. Approximately 70K unique individuals completed the survey, on average 17,468 individuals a week in the UK (Fancourt et al., 2021). Specifically, participants were asked to indicate if over the previous week they felt worried and/or stressed about any of the following factors: (i) catching Covid-19, (ii) becoming seriously ill from Covid-19, (iii) finance, (iv) unemployment and (v) getting food. We computed the mean proportion of individuals who reported stress or worry over these factors. We conducted separate linear models, each predicting on a weekly basis either the Google search volume index in the UK of *"How"* searches from stress and from UK COVID-19 confinement scores. The inclusion of the latter enabled us to disentangle the effects of stress on web searches from the effects of confinement due to restrictions placed by the Government. Covid-19 related confinement data for each week in the UK was obtained from the Oxford University COVID-19 Government Response Tracker (Webster et al., 2021). The data includes ordinal variables coded by severity/intensity of confinement, on a daily basis (from January 1st, 2020 and March 21st, 2021), due to the following: (i) school and university closures, (ii) workplace closures, (iii) public event cancelations, (iv) restrictions on gatherings, (v) public transport restrictions, (vi) stay at home requirements, (vii) restrictions on domestic travel, and (viii) restrictions on international travel; **see Table 4.3 for coding**. To obtain weekly values, we computed weekly averages of the daily ratings, where all eight variables were normalised to range between 0 and 1 and averaged together.

Importantly, to account for simple temporal trends we removed linear trend (e.g., Lampos et al., 2021; Bakker et al., 2016) from the dependent variables and predictor variables (stress scores and COVID-19 related confinement), using the detrend function in the 'pracma' R package. The detrended dependent and predictor variables were then Z-scored before being entered in the linear models.

The linear model predicting *"How"* questions from stress levels and COVID-19 related confinement scores, revealed that both high stress ($\beta = 0.182 \pm 0.074$ (SE), $t(49) = 2.464$, $p = 0.017$) and greater COVID-19 related confinement ($\beta = 0.797 \pm$

0.074 (SE), t(49) = 10.812, p < 0.0001; **Figure 4.3a**) predicted proportion of *"How"* searches. In other words, the relationship between stress levels and *"How"* searches cannot be solely explained by increased restrictions during the pandemic, as our model controls for COVID-19 related confinement.

We then conducted the same linear model as above, but this time predicting the Valence Index. The Valence of searches was not predicted by either variable (stress: β = -0.144 ± 0.147 (SE), t(49) = -0.980, p = 0.332, COVID-19 related confinement: β = 0.216 ± 0.147 (SE), t(49) = 1.470, p = 0.148; **Figure 4.3b**).

Next, to examine if the relationship between stress and *"How"* searches was specific, or rather reflected a general tendency to ask more question when stressed, we conducted a third linear model relating question asking to stress levels controlling for confinement. In particular, we extracted the Google search volume index for all other common question-words (i.e., What, Which, Who, Where, Why, When, and Whose; Z-scored and detrended) and then averaged these together. We then predicted the average Google search volume index of all other common question-words from stress scores and COVID-19 related confinement. Importantly, the proportion of other common questions asked was selectively predicted by COVID-19 related confinement (β = 0.701 ± 0.104 (SE), t(49) = 6.744, p < 0.0001), but not stress levels (β = 0.067 ± 0.104 (SE), t(49) = 0.645, p = 0.522) (**Figure 4.3c**).

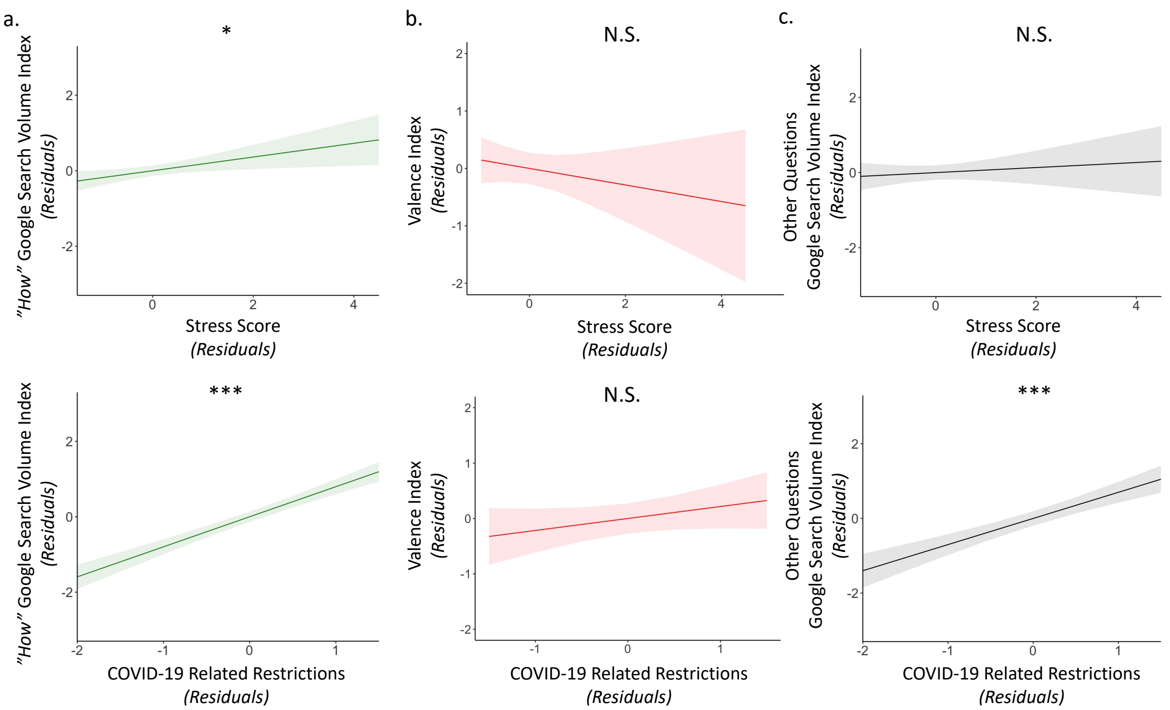

Figure 4.3. Self-reported stress is selectively associated with an increase in *"How"* searches. Stress level in the UK was associated with **(a)** the UK Google search volume index of *"How"* searches (detrended and Z-scored), but not with **(b)** the Valence Index [0 (most negative valenced) and 100 (most positive valenced)] (detrended and Z-scored), nor with **(c)** the mean UK Google search volume index of

other questions (i.e., What, Which, Who, Where, Why, When, and Whose; detrended and Z-scored). UK COVID-19 related confinement score (detrended and Z-scored; bottom panel) was associated with both **(a**) the UK Google search volume index of *"How"* searches and **(c)** the mean UK Google search volume index of other questions, but not **(b)** the Valence Index. Stress levels and COVID-19 confinement scores (all detrended and Z-scored) were entered in the same models, controlling for each other. The X and Y values are the residuals (regressing out the respective control variable). The fine line represents the confidence interval. ***P < 0.001, *P < 0.05, N.S. = not significant (two-sided).

Thus far, we have shown that the relative volume of searches that can direct action is related to stress levels. Next, we wanted to test the predictive validity of this simple model. Specifically, we used our stress score to predict the proportion of *"How"* searches using a leave one out analysis. To account for a simple temporal trend, we first removed the linear trend from the dependent variable (*"How"* questions) and the predictor variable (stress levels). The detrended predictor variables were then Z-scored before being entered in the simple linear model. The simple model was then run on all the data save for one time point which was held out from the analysis. We then used the regression beta to predict the proportion of *"How"* searches of the left-out time point. This process was repeated so that each week's proportion of *"How"* searches was estimated from the simple model parameters generated without using that week to fit the data. The actual proportion of *"How"* searches of a week (data) and the predicted proportion of *"How"* searches (estimation) were then correlated and also compared using a paired sample t-test. We observed a correlation between the predicted proportion of *"How"* searches (estimate) and the actual proportion of *"How"* searches (data) ($r(50) = 0.366$, $p = 0.007$). The means of the two sets of values were not significantly different from one another ($t = -0.041$, $p = 0.967$). This analysis suggests that stress levels in a population is a good predictor of the proportion of *"How"* searches during the pandemic.

Next, we tested whether stress was better predicted by *"How"* Google searches than searches for specific content terms (i.e., *'stress'*, *'anxiety'*, *'mental health'* and *'psychiatrist'*), which are often used in attempt to predict population mental state. Thus, we ran multiple linear models to predict stress from each term separately. Once again, the dependent and predictor variables were first detrended and then Z-scored. The strongest association was seen by *"How"* question volume and self-reported stress scores ($\beta = 0.439 \pm 0.127$ (SE), $t(50) = 3.435$, $p = 0.001$, R2.= 0.437), followed by the Google search index "*stress*" ($\beta = 0.322 \pm 0.134$ (SE), $t(50) = 2.403$, $p = 0.02$, R2.= 0.322) and *'psychiatrist'* ($\beta = -0.348 \pm 0.131$ (SE), $t(50) = -2.887$, $p = 0.006$, R2.= -0.378); all other predictors were not significant (p's >= 0.594). Note that the relationship between the volume of *'psychiatrist'* searches and self-reported stress was inverse. This may be due to a decreased access to in-person psychiatrist sessions during lockdown, which would correspond to increasing stress.

In the US, we did not have access to measurements of population stress levels. However, we did have access to COVID-19 related confinement data which enabled us to examine the relationship between web searches and residential confinement in the US, when stress levels are not controlled for. We observed that increased COVID-19 related confinement was related to greater Google search volume index "*How*" ($\beta = 0.384 \pm 0.129$ (SE), $t(50) = 2.940$, $p = 0.005$, R2.= 0.384) and to more negatively valenced searches ($\beta = -0.316 \pm 0.134$ (SE), $t(50) = -2.354$, $p = 0.023$, R2.= -0.316). We did not observe a relationship between COVID-19 related confinement and the average Google search volume index of the other question-words ($\beta = -0.142 \pm 0.140$ (SE), $t(50) = -1.015$, $p = 0.315$, R2.= -0.142). Note, all results presented above remain when not detrending.

Study 3

Personal stressful events are associated with an increase in *"How"* and negative valenced searches. To assess whether stressful events influence the propensity to ask "*How*" and the valence of questions in other situations, we ran a third study. First, we asked participants to recall in detail, and write about, either a stressful past event (e.g., *"I had a deadline at work and didn't know if I was going to meet it."*; stress condition; n = 99, age = 39.45, SD = 14.74; Females = 78.8%, Males = 19.2%, Other = 2.0%) or a relaxing and happy past event (e.g., *"My holiday in [retracted] with my aunt, cousin and her children. The weather was great, nice and warm. We were staying at a resort on the beach."*); control condition; n = 93, age = 38.46, SD = 11.91; Females = 76.3%, Males = 22.6%, Other = 1.1%). Participants reported their stress level on a scale ranging from *very calm* (-50) to *very stressed* (+50) before and after recalling the event (**see Methods for details**).

A 2 (condition: stress, control) by 2 (time: pre-induction, post induction) ANOVA on self-reported stress revealed a significant interaction ($F(1, 190) = 43.074$, $p < 0.001$, partial eta squared = 0.181). Post-hoc pair-wise t-tests revealed that the interaction was characterised by participants in the stress condition reporting higher stress post induction (M = -3.59, SD = 24.30) compared to pre-induction (M = -12.57, SD = 24.94, $t(98) = 5.158$, $p < 0.001$, Cohen's d = 0.518), in contrast, in the control condition, participants mood significantly increased post induction compared to pre-induction (Post induction: M = -18.80, SD = 27.34, pre-induction: M = -13.38, SD = 27.79, $t(92) = -4.028$, $p < 0.001$, Cohen's d = -0.418). Importantly, participants in the stress condition reported greater stress post induction relative to controls (Stress condition: M = -3.59, SD = 24.30; Control condition: M = -18.80, SD = 27.34, $t(190) = 2.581$, $p < 0.001$, Cohen's d = 0.589), with no difference pre-induction (Stress condition: M = -12.57, SD = 24.94; Control condition: M = -13.38, SD = 27.79, $t(190) = 1.712$, $p = 0.832$, Cohen's d = 0.031; **see Figure 4.4a**).

Next, participants were asked to enter two searches they could have entered to Google during the time. For each participant, we counted the number of questions that began with *"How"* (i.e., 0,1,2) and also assessed the average valence of participants' two searches, by applying the pre-trained98 model distilbert-base-uncased-finetuned-sst-2- english (HuggingFace, 2022). Results show that participants in the stress condition asked significantly more *"How"* questions (M = 0.82, SD = 0.77) than those in the control condition (M = 0.16, SD = 0.28, $t(124.58) = 94.154$, $p < 0.001$;

Cohen's d = 1.110; **see Figure 4.4b**). In addition, participants in the stress condition also submitted significantly more negative searches (M = 40.55, SD = 36.51) than those in the control condition (M = 53.97, SD = 36.11, t(190) = -2.558, p = 0.011; Cohen's d = -0.369; **see Figure 4.4c**). This study strengthens the conclusion of Study 2, that stress inducing events are associated with asking more *"How"* questions and negatively valenced question by generalising the finding of Study 2 to other stressful events and providing support for the conclusion in a controlled setting. As shown in **Table 4.4**, there was no increase in the use of other questions words under stress relative to control.

Table 4.4. An exploratory analysis showed that there was no increase in the use of any other question words under stress relative to control. The use of "*What"* showed a decrease under stress relative to control.

Question Type	***Mean–Difference (Stress - Control)***	***t-test Stats***
What	-0.43	$t(144.65) = -5.179$, $p < 0.001$
When	-0.03	$t(152.34) = -1.221$, $p = 0.224$
Which	0.01	$t(190) = 0.969$, $p = 0.334$
Why	0.01	$t(190) = 0.244$, $p = 0.807$
Where	-0.08	$t(141.80) = -1.855$, $p = 0.066$
Who	-0.02	$t(143.47) = -1.055$, $p = 0.293$

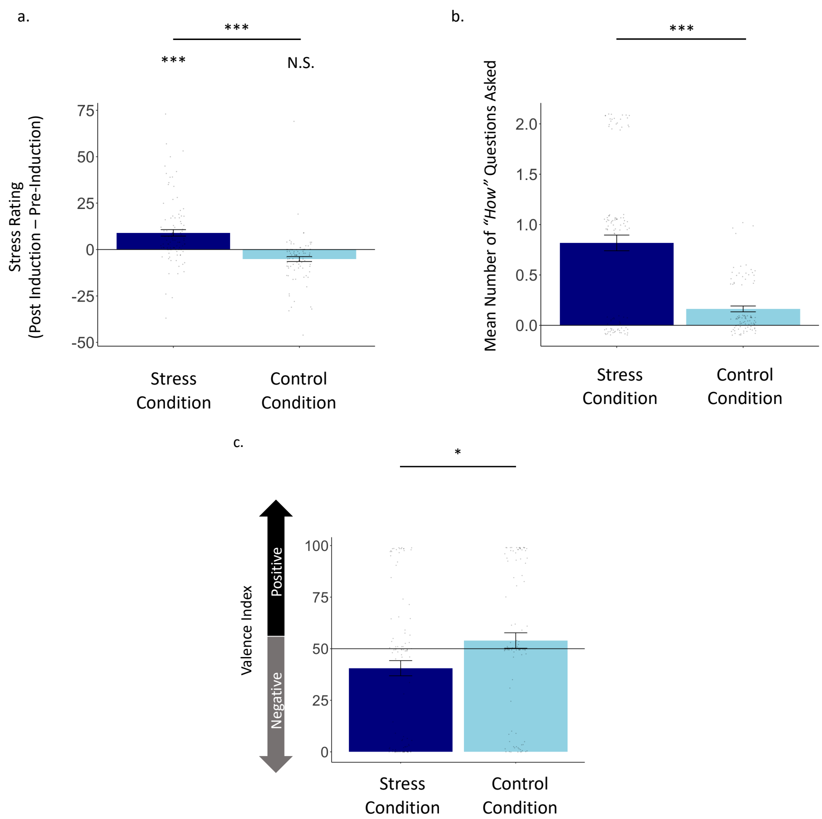

Figure 4.4. Personal stressful events alter high-level characteristics of searches. (a) Participants were asked to recall in detail, and write about, either a stressful past event (i.e., stress condition) or a relaxing past event (i.e., control condition). Participants reported their stress level on a scale ranging from *very calm* (-50) to *very stressed* (+50) before and after recalling the event. Plotted on the y-axis is participant's post induction stress rating minus their pre-induction stress rating for the stress condition (dark blue) and control condition (light blue). Participants' stress scores increased post stress induction compared to pre-induction for the stress condition but not the control condition. **(b)** The mean number of *"How"* questions asked in the stress condition (x-axis; dark blue) was greater than in the control condition (x-axis; light blue). **(c)** In the stress condition (x-axis; dark blue) participants asked more negative questions than in the control condition (x-axis; light blue). Individual scores are shown as dots. Error bars = standard error (SEM). *** = P < 0.001, * = P < 0.05, N.S. = not significant (two-sided).

4.4 Discussion

The global pandemic generated a new set of practical and mental challenges. To overcome these challenges, people turned to technology. On average, people spent almost 7-hours a day online in 2020, up 7.3% from the previous year (Kemp, 2021). A large fraction of this time was dedicated to searching for and consuming information (Kemp, 2021). Here, we examined how the high-level features of searches submitted to Google changed in response to the pandemic and private aversive life events and how such changes relate to stress.

In particular, we were interested if under stress people would seek more information that could guide their actions – a behaviour that could facilitate the process of adapting to a stressful event. Study 1 revealed that participants specifically use the word *"How"* when they want information to guide their action. Thus, in Study 2 we examined the frequency of Google searches that included the word *"How"*. We found that both in the UK and US, the proportion of searches that included the word *"How"* were greater during the year following the declaration of *"National Emergency"* than in the years prior. It is important to emphasise that any change in the proportion of *"How"* searches cannot simply be explained by a general increase in number of Google searches, as the former are calculated as proportion of the latter. Neither can it be explained by temporal linear trends, as the data was detrended. The rise in *"How"* searches may reflect an adaptive human tendency to ask questions that can facilitate rapid adjustment to new and potentially aversive environments. We were also interested in changes to the valence of Google searches during the pandemic relative to before. We observed that the most popular searches in the UK and US were more negatively valenced during the year following the declaration of *"National Emergency"* than in the years prior. This aligns with the notion that people may search for information that aligns with their emotional state.

This strategy of seeking instrumental information may be particularly effective when individuals have agency regarding the event causing stressed, as research indicates that having a sense of control can reduce stress (Bandura, 1997; Frazier, Berman, & Steward, 2001). For instance, studies have shown that when individuals perceive they have control over a situation, they experience lower stress levels (Bandura, 1997; Frazier, Berman, & Steward, 2001). However, in scenarios where individuals have little or no control over the outcome, such as in experiments where participants receive electric shocks, this strategy is less likely to be effective. In these cases, the lack of agency can exacerbate stress rather than alleviate it (Geer, Davison, & Gatchel, 1970). To test our prediction that people are more likely to ask "*How*" questions during stressful times when they have some control over the outcome, we designed a study presenting participants with controllable and uncontrollable scenarios. For instance, participants faced scenarios such as an immediate, impromptu public speaking invitation (uncontrollable) versus a speaking event scheduled a week in advance (controllable). We predict that in the controllable scenario, participants will ask more "*How*" questions (e.g., *How to design a good presentation?),* while in the uncontrollable scenario, participants might ask other types of questions, such as "*What now?*" etc. Relatedly, it is important to assess how other emotional states such as happiness and anger influence the types of information we seek.

Weekly fluctuations in the proportion of *"How"* questions submitted to Google during the pandemic was positively associated with weekly fluctuations in the proportion of individuals who reported experiencing COVID-related stress in the UK in a sample of over 17K residences. This association could not be attributed to COVID-related confinement, as this factor was controlled for in the model. Furthermore, the relationship was specific to *"How"* searches, and did not generalise to general question asking.

Markedly, we show that the frequency of *"How"* searches predicted COVID-related stress better than the frequency of searches that include stress related content (i.e., '*anxiety*', '*stress*', and '*mental health*'). This raises the novel idea that tracking the frequency of *"How"* searches may predict population-level stress beyond the time of a pandemic and perhaps predict an individual's stress level. If affirmative, quantifying such search features can prove extremely valuable for monitoring stress on both an individual and populating level. The findings are in accord with a recent study suggesting that changes to search during the pandemic reflect a change in people's needs (Suh et al., 2021).

Consistent with the above, in Study 3 we show that people are also more likely to ask "*How*" questions, and more negatively valenced questions, in response to a personal stressful life event than a control one. In particular, participants were instructed to recall in detail a stressful life event or a calming event. The former, but not the latter, increased self-reported stress. They were then asked which questions they could have submitted at the time. Recalling stressful events was associated with significantly more *"How"* questions and negative questions than control events. This suggests that the relationship between asking *"How",* and stress is not specific to the pandemic but extends to stressful events in general.

Our investigation was guided by previously identified factors that motivate information-seeking (for review see Sharot & Sunstein, 2020). In particular, studies show that people seek information more when it is useful in guiding action (Kelly & Sharot, 2021; Stigler,1961; Kobayashi & Hsu, 2019; Cogliati-Dezza et al., 2022). The current results suggest that this motive is ‘up weighted’ when experiencing stress, perhaps because the need to select adaptive actions is heightened under such circumstances. While past studies also show people prefer to seek good news over bad (for review see Sharot & Sunstein, 2020; Kelly & Sharot, 2021; Karlsson et al., 2009; Golman et al., 2017; Persoskie et al., 2014; Vellani et al., 2020; Charpentier at al., 2018; Lerman et al., 1998; Kobayashi et al., 2019; Cogliati-Dezza et al., 2022) we find that under stress web searches were in fact more negative. This could simply be due to a significant proportion of searches relating to the stress itself (i.e., illness in Study 2 and a range of personal aversive events in Study 3). We did not, however, observe a significant association between valence of web searches and self-reported stress levels. It is interesting to note that studies examining information-*sharing* (e.g., tweets) rather than information-*seeking* have revealed small but significant associations between mental state and the valence of information *shared* (De Choudhury et al., 2013; Kelley & Gillan, 2022; Eichstaedt et al., 2018). Because a significantly smaller slice of the population regularly shares than seeks information, understanding how information-seeking relates to mental health is crucial and may diverge from patterns observed for information-*sharing*.

Together, the findings show that in the face of a novel stressful situation and high stress people search for information that can help guide action. While in the past such information may have been sought directly from other people, with the development of the internet, individuals are now able to turn to the web for answers. This ability may have contributed to the high resilience and quick adaptation observed in response to the pandemic (Aknin et al., 2021; Globig et al., 2022; Daly et al., 2021).

4.5 References

Aknin, L., Neve, J.-E. De, Dunn, E., Fancourt, D., Goldberg, E., Helliwell, J., Jones, S. P., Karam, E., Layard, R., Lyubomirsky, S., Rzepa, A., Saxena, S., Thornton, E., VanderWeele, T., Whillans, A., zaki, jamil, Caman, O. K., & Amour, Y. Ben. (2021). Mental Health During the First Year of the COVID-19 Pandemic: A Review and Recommendations for Moving Forward. https://doi.org/10.31234/OSF.IO/ZW93G

Ayers, J. W., Althouse, B. M., Allem, J. P., Childers, M. A., Zafar, W., Latkin, C., ... & Brownstein, J. S. (2012). Novel surveillance of psychological distress during the great recession. Journal of Affective Disorders, 142(1–3), 323–330. https://doi.org/10.1016/J.JAD.2012.05.005

Ayers, J. W., Poliak, A., Johnson, D. C., Leas, E. C., Dredze, M., Caputi, T., & Nobles, A. L. (2021). Suicide-Related Internet Searches During the Early Stages of the COVID-19 Pandemic in the US. JAMA Network Open, 4(1). https://doi.org/10.1001/JAMANETWORKOPEN.2020.34261

Bakker, K. M., Martinez-Bakker, M. E., Helm, B., & Stevenson, T. J. (2016). Digital epidemiology reveals global childhood disease seasonality and the effects of immunization. Proceedings of the National Academy of Sciences, 113(24), 6689-6694.

Barros, J. M., Melia, R., Francis, K., Bogue, J., O'Sullivan, M., Young, K., ... & Duggan, J (2019). The Validity of Google Trends Search Volumes for Behavioral Forecasting of National Suicide Rates in Ireland. International Journal of Environmental Research and Public Health, 16(17). https://doi.org/10.3390/IJERPH16173201

Bar-Haim, Y., Lamy, D., Pergamin, L., Bakermans-Kranenburg, M. J., & Van Ijzendoorn, M. H. (2007). Threat-related attentional bias in anxious and nonanxious individuals: a meta-analytic study. Psychological bulletin, 133(1), 1.

Chae, J. (2016). Who avoids cancer information? Examining a psychological process leading to cancer information avoidance. Journal of health communication, 21(7), 837-844.

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 of the United States of America. https://doi.org/10.1073/pnas.1800547115

Cisler, J. M., & Koster, E. H. (2010). Mechanisms of attentional biases towards threat in anxiety disorders: An integrative review. Clinical psychology review, 30(2), 203-216.

Daly, M., & Robinson, E. (2021). Psychological distress and adaptation to the COVID-19 crisis in the United States. Journal of psychiatric research, 136, 603-609.

De Choudhury, M., Gamon, M., Counts, S., & Horvitz, E. (2013). Predicting depression via social media. In Seventh international AAAI conference on weblogs and social media.

Dezza, I. C., Maher, C., & Sharot, T. (2022). People adaptively use information to improve their internal states and external outcomes. Cognition, 228, 105224.

Drouin, M., McDaniel, B. T., Pater, J., & Toscos, T. (2020). How parents and their children used social media and technology at the beginning of the COVID-19 pandemic and associations with anxiety. Cyberpsychology, Behavior, and Social Networking, 23(11), 727-736.

Ebrahim, A. H., Saif, Z. Q., Buheji, M., AlBasri, N., Al-Husaini, F. A., & Jahrami, H. (2020). COVID-19 information-seeking behavior and anxiety symptoms among parents. OSP Journal of Health Care and Medicine, 1(1), 1-9.

Eichstaedt, J. C., Smith, R. J., Merchant, R. M., Ungar, L. H., Crutchley, P., Preoţiuc-Pietro, D., ... & Schwartz, H. A. (2018). Facebook language predicts depression in medical records. Proceedings of the National Academy of Sciences, 115(44), 11203-11208.

Fancourt, D., Bu, F., Mak, H. W., & Steptoe, A. (2021). COVID-19 social study. Results release, 25.

Fancourt, D., Steptoe, A., & Bu, F. (2021). Trajectories of anxiety and depressive symptoms during

enforced isolation due to COVID-19 in England: a longitudinal observational study. The Lancet Psychiatry, 8(2), 141-149.

Finlay-Jones, R., & Brown, G. W. (1981). Types of Stressful Life Event and the Onset of Anxiety and Depressive Disorders. Psychological Medicine, 11(4), 803–815. https://doi.org/10.1017/S0033291700041301

Francis, J. L., Moitra, E., Dyck, I., & Keller, M. B. (2012). The impact of stressful life events on relapse of generalized anxiety disorder. Depression and Anxiety, 29(5), 386–391. https://doi.org/10.1002/da.20919

Globig, L. K., Blain, B., & Sharot, T. (2022). Perceptions of personal and public risk: Dissociable effects on behavior and well-being. Journal of Risk and Uncertainty, 1-22.

Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information avoidance. In Journal of Economic Literature. https://doi.org/10.1257/jel.20151245

Gunnell, D., Derges, J., Chang, S. Sen, & Biddle, L. (2015). Searching for suicide methods. Crisis, 36(5), 325–331. https://doi.org/10.1027/0227-5910/A000326

Hirshleifer, J., & Riley, J. G. (1979). The Analytics of Uncertainty and Information- An Expository Survey. Source Journal of Economic Literature Journal of Economic Literature.

Hoerger, M., Alonzi, S., Perry, L. M., Voss, H. M., Easwar, S., & Gerhart, J. I. (2020). Impact of the COVID-19 pandemic on mental health: Real-time surveillance using google trends. Psychological Trauma: Theory, Research, Practice, and Policy, 12(6), 1–2. https://doi.org/10.1037/TRA0000872

HuggingFace. 2022. DistilBERT base uncased finetuned SST-2. https://huggingface.co/distilbert-base-uncased-finetuned-sst-2-english. Accessed: 2022- 07-13.

Karlsson, N., Loewenstein, G., & Seppi, D. (2009). The ostrich effect: Selective attention to information. Journal of Risk and Uncertainty. https://doi.org/10.1007/s11166-009-9060-6

Kash, K. M., Ortega-Verdejo, K., Dabney, M. K., Holland, J. C., Miller, D. G., & Osborne, M. P. (2000, June). Psychosocial aspects of cancer genetics: women at high risk for breast and ovarian cancer. In Seminars in Surgical Oncology (Vol. 18, No. 4, pp. 333-338). New York: John Wiley & Sons, Inc.

Kelly, C.A., Sharot, (2021) T. Individual differences in information-seeking. Nat Commun 12, 7062. https://doi.org/10.1038/s41467-021-27046-5

Kelley, S. W., & Gillan, C. M. (2022). Using language in social media posts to study the network dynamics of depression longitudinally. Nature communications, 13(1), 1-11.

Kemp. S. (2021). Digital 2021: Global Overview Report — DataReportal – Global Digital Insights. https://datareportal.com/reports/digital-2021-global-overview-report

Kobayashi, K., & Hsu, M. (2019). Common neural code for reward and information value. Proceedings of the National Academy of Sciences, 116(26), 13061-13066.
Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M. & Gottlieb, J. Diverse motives for human curiosity. Nat. Hum. Behav. 3, 587–595 (2019).
Knipe, D., Evans, H., Marchant, A., Gunnell, D., & John, A. (2020). Mapping population mental health concerns related to COVID-19 and the consequences of physical distancing: a Google trends analysis. Wellcome Open
Lampos, V., Majumder, M. S., Yom-Tov, E., Edelstein, M., Moura, S., Hamada, Y., ... & Cox, I. J. (2021).
Tracking COVID-19 using online search. NPJ digital medicine, 4(1), 1-11. Research 2020 5:82, 5, 82. https://doi.org/10.12688/wellcomeopenres.15870.2
Lerman, C. et al. (1998). What you don't know can hurt you: adverse psychologic effects in members of BRCA1-linked and BRCA2-linked families who decline genetic testing. J. Clin. Oncol. 16, 1650–1654.
Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. Psychological bulletin, 116(1), 75.
Loosen, A. M., Skvortsova, V., & Hauser, T. U. (2021). Obsessive–compulsive symptoms and information seeking during the Covid-19 pandemic. Translational psychiatry, 11(1), 309.
MacLeod, C., & Mathews, A. (1988). Anxiety and the allocation of attention to threat. The Quarterly journal of experimental psychology, 40(4), 653-670.
McCarthy, M. J. (2010). Internet Monitoring of Suicide Risk in the Population. Journal of Affective Disorders, 122(3), 277. https://doi.org/10.1016/J.JAD.2009.08.015
McLaughlin, K. A., Conron, K. J., Koenen, K. C., & Gilman, S. E. (2010). Childhood adversity, adult stressful life events, and risk of past-year psychiatric disorder: A test of the stress sensitization hypothesis in a population-based sample of adults. Psychological Medicine, 40(10), 1647–1658. https://doi.org/10.1017/S0033291709992121
Miloyan, B., Joseph Bienvenu, O., Brilot, B., & Eaton, W. W. (2018). Adverse life events and the onset of anxiety disorders. Psychiatry Research, 259, 488–492. https://doi.org/10.1016/j.psychres.2017.11.027
Misiak, B., Szcześniak, D., Koczanowicz, L., & Rymaszewska, J. (2020). The COVID-19 outbreak and Google searches: Is it really the time to worry about global mental health? Brain, Behavior, and Immunity, 87, 126. https://doi.org/10.1016/J.BBI.2020.04.083
Persoskie, A., Ferrer, R. A. & Klein, W. M. (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.
Rana, U. (2020). Elderly suicides in India: An emerging concern during COVID-19 pandemic. International Psychogeriatrics, 32(10), 1251–1252. https://doi.org/10.1017/S1041610220001052
Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. Nature Human Behaviour. https://doi.org/10.1038/s41562-019-0793-1
Sinyor, M., Spittal, M. J., & Niederkrotenthaler, T. (2020). Changes in Suicide and Resilience-related Google
Searches during the Early Stages of the COVID-19 Pandemic. Canadian Journal of Psychiatry. Revue Canadienne de Psychiatrie, 65(10), 741–743. https://doi.org/10.1177/0706743720933426

Stigler, G. J. (1961). The Economics of Information. Journal of Political Economy. https://doi.org/10.1086/258464

Suh, J., Horvitz, E., White, R. W., & Althoff, T. (2021). Population-scale study of human needs during the covid-19 pandemic: Analysis and implications. In Proceedings of the 14th ACM International Conference on Web Search and Data Mining (pp. 4-12).

Suls, J., & Mullen, B. (1981). Life Events, Perceived Control and Illness: The Role of Uncertainty. Journal of Human Stress, 7(2), 30–34. https://doi.org/10.1080/0097840X.1981.9936824

Sueki, H. (2011). Does the volume of Internet searches using suicide-related search terms influence the suicide death rate: data from 2004 to 2009 in Japan. Psychiatry and Clinical Neurosciences, 65(4), 392–394. https://doi.org/10.1111/J.1440-1819.2011.02216.X

Tran, U. S., Andel, R., Niederkrotenthaler, T., Till, B., Ajdacic-Gross, V., & Voracek, M. (2017). Low validity of Google Trends for behavioral forecasting of national suicide rates. PLOS ONE, 12(8), e0183149. https://doi.org/10.1371/JOURNAL.PONE.0183149

Vellani, V., de Vries, L. P., Gaule, A., & Sharot, T. (2020). A selective effect of dopamine on information-seeking. Elife, 9, e59152.

WHO. (2021). WHO Coronavirus (COVID-19) Dashboard With Vaccination Data. https://covid19.who.int/

Yang, A. C., Huang, N. E., Peng, C.-K., & Tsai, S.-J. (2010). Do Seasons Have an Influence on the Incidence of Depression? The Use of an Internet Search Engine Query Data as a Proxy of Human Affect. PLOS ONE, 5(10), e13728. https://doi.org/10.1371/JOURNAL.PONE.0013728

Zhang, D., Hu, M., & Ji, Q. (2020). Financial markets under the global pandemic of COVID-19. Finance Research Letters, 36. https://doi.org/10.1016/j.frl.2020.101528

Frazier, P., Berman, M., & Steward, J. (2001). Perceived control and posttraumatic stress: A temporal model. Applied and Preventive Psychology, 10(3), 207-223.

Bandura, A. (2013). Self-Efficacy: The Foundation of Agency1. In Control of human behavior, mental processes, and consciousness (pp. 16-30). Psychology Press.

Geer, J. H., Davison, G. C., & Gatchel, R. I. (1970). Reduction of stress in humans through nonveridical perceived control of aversive stimulation. Journal of Personality and Social Psychology, 16(4), 73

Chapter 5

A Tool to Facilitate Web-Browsing

5.1 Overview

Approximately eight billion search engine queries are submitted daily by individuals who seek to gain knowledge and make informed decisions (Kemp, 2022). However, search results are shaped by opaque algorithms that do not necessarily align with users' goals (Rainie, Lee & Anderson, 2017). Consequently, individuals dedicate countless hours to absorbing information that may not yield practical benefits, and in some cases, may have a detrimental effect on their well-being (Kelly & Sharot, 2023). For example, by consuming negatively valanced information that is not informative or helpful.

To address this problem we developed a tool designed to empower users to navigate the web in a way that may improve their decision making, mental health, and understanding. Much like how people use nutritional labels to learn about the nutritional value of food before it enters their body (e.g., calories, fat content etc.), the tool provides 'content labels' for available webpages in a search engine results page that a user can inspect before consuming information.

In particular, the software (in the form of a Google Chrome plugin) informs users of three properties that can guide information-consumption decisions: (i) actionability (the ability of text on a webpage to guide action, on average); (ii) ability of text on a web page to enhance understanding, on average; (iii) sentiment (e.g., how positive or negative the text on a webpage is). These three properties were selected based on empirical research that indicates that people's key motives for seeking information is to (i) guide their actions and decisions (Kelly & Sharot, 2021; Cogliati Dezza et al., 2022; Stigler, 1961; for review Sharot & Sunstein, 2020), (ii) improve comprehension (Cogliati Dezza et al., 2022; Sharot & Sunstein, 2020) and (iii) improve affect (Kelly & Sharot, 2021; Sharot & Sunstein, 2020; Charpentier et al., 2018; Loewenstein, 1994; Caplin & Leahy, 2001; Kőszegi, 2010; Golman et al., 2017).

The plugin provides scores visible in a Google search results (**see Figure 5.1**) about the above three factors [which we respectively called 'Actionability', 'Knowledge', and 'Emotion'] of text found on webpages. Users can use these scores to improve their web-browsing experience, such that the information they consume better aligns with their goals. For instance, individuals seeking practical advice such as *"I just lost my job"* may prioritise information with a high 'Actionability' value, while those looking to deepen their understanding of a topic, for example *"who is the most famous pharaoh"* might prioritise webpages with a high 'Knowledge' score.

Indeed, different individuals may prioritise some of these scores over others (Kelly & Sharot, 2021). For example, some people may be driven more to seek information that can help them make better decisions, while other may be primarily driven to seek information that helps them understand the world better. The importance of these motives can vary as a function of a person's state (e.g., stress vs. relaxed state etc.) and domain (for example, in the domain of health 'Actionability' of information may be especially important; Kelly & Sharot, 2021).

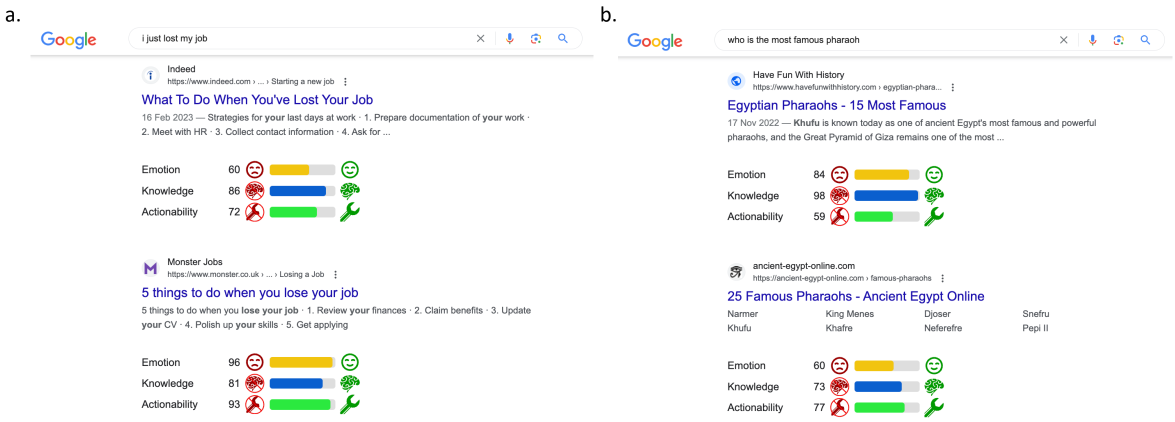

Figure 5.1. Presentation of Scores. The figure presents the scores for webpages obtained from the Google search engine for two different queries: **(a)** *"I just lost my job"* and **(b)** *"who is the most famous pharaoh"*. The *Emotion* (yellow), *Knowledge* (blue), and *Actionability* (green) scores are computed for each webpage listed in the Google search results using the process described in the **Tool Development** section below. The user is then presented with these scores alongside each webpage. This feature enables the user to make informed decisions about which webpage to visit, which can improve their web-browsing experience. **(a)** For the first search term, *"I just lost my job"*, the second result offers the most positive *Emotion* and also the highest *Actionability* score. these two metrics might be particularly relevant for the specific online search objective – to obtain information that can address a challenging situation, while maintaining a positive tone. **(b)** For the second search term, *"who is the most famous pharaoh"*, the top result showcased the highest Knowledge score, potentially aligning best with a user's objective of enhancing their understanding of the topic.

The nature of how webpages are interpreted and rated is of course subjective. For instance, a webpage that is perceived as positive by one person may be perceived negatively by another. Yet, as the results detailed below demonstrate there is nonetheless high agreement across users ***on average*** regarding the valence, actionability and potential knowledge enhancement of webpages. This suggests that despite subjectivity and individual differences, it is possible to effectively capture a shared perception that is relevant to many users and can be leveraged. Just as mean ratings of products (books, movies, items) are often helpful despite their subjective nature, 'on average' scores of websites can be valuable in guiding users' online information consumption, allowing them to engage with information that aligns with their goals and preferences.

5.2 Tool Development

To measure and present the scores of interest along the Google search results we applied the following method (**see Figure 5.2**):

<u>Webpage Retrieval & Parsing:</u> For each Google search submitted by users, we extracted the base HyperText Markup Language (HTML) source code from each of the web pages and parsed the code using the Python package 'beautifulsoup4' (Richardson, 2007). We then extracted the paragraph text from each of those web pages.

<u>Emotion Scoring:</u> To quantify Emotion of webpages on a Google search results page, we chose to employ the VADER (Valence Aware Dictionary and sEntiment Reasoner) sentiment analysis tool. This is a lexicon and rule-based tool that was developed by Hutto and Gilbert (2014) and is particularly efficient at capturing sentiments expressed in media content.

The VADER sentiment analysis tool assigns a sentiment score, known as valence, to each word in the text. This valence score typically falls within the range of -4 (extremely negative) to +4 (extremely positive), with 0 signifying neutral sentiment. After these scores are assigned, the compound score is then calculated.

The calculation of this compound score follows these steps:

1. The text from each webpage is parsed into individual words or tokens.
2. Each token is matched against the VADER lexicon and when a match is found, the corresponding valence score is recorded.
3. The initial valence score is adjusted based on contextual factors. For example, sentiment scores can be altered if the word is preceded by a negation word like "not" or "isn't", and can be influenced by punctuation, capitalisation, degree modifiers, and contrastive conjunctions such as "but".
4. The adjusted sentiment scores for all words are then summed up to get a total sentiment score for the text of the webpage.

Despite the individual words having a score range from -4 to +4, the ultimate compound score is adjusted to fit within a -1 to +1 scale. This makes the Vader sentiment analysis tool a fitting choice for our needs, given its straightforwardness in interpretation. This is achieved through the following formula:

$$Compound\ score = \frac{\sum Valence\ scores\ of\ each\ word}{\sqrt{\sum Valence\ scores\ of\ each\ word^2 + 15}} \quad (1)$$

This "15" in the formula is a constant that was empirically determined through testing by Hutto and Gilbert (2014) to provide the most accurate normalisation. Its role is essential in ensuring that the compound score falls within the -1 to +1 range, irrespective of the length of the text or the individual valence scores.

Essentially, the compound score offers a holistic sentiment metric for each webpage. A higher score implies a more positive sentiment, a lower score suggests a negative sentiment, and a score around zero denotes a neutral sentiment. This methodology respects the context and the subtle language nuances that often impact sentiment, thereby providing a more accurate and comprehensive measure of overall sentiment.

Actionability & Knowledge Scoring:

1. *Participant Ratings and Label Collection*

   1000 participants (Age = 39.07 (SD = 13.12), Female = 46.2%, Male = 53.3%, Other = 0.5%) were recruited via Prolific's online recruitment platform (www.prolific.co) to browse and rate 5 webpages each on two dimensions: Actionability and Knowledge (total ratings per dimension = ~ 5,000). Participants were paid at a rate of £9.00 per hour for their participation. Ethical approval was provided by the Research Ethics Committee at UCL, and all participants gave their informed consent to participate.

   Actionability was defined as the extent to which the information on the webpage could help guide actions and/or decisions (i.e., *"Could the information on the webpage help guide actions and/or decisions?"*). Knowledge was defined as the degree to which the information on the webpage increased the participant's understanding of the topic (i.e., *"Does the information on the webpage increase your understanding of the topic?"*). Both dimensions were rated on a 6-point scale, with 1 representing the lowest level of Actionability and Knowledge, and 6 representing the highest level.

2. *Model Training and Evaluation*

   The model training and evaluation process was performed following a structured sequence of steps. First, given that the Actionability and Knowledge scores ranged from 1 (low) to 6 (high), we set a binary threshold. The optimal overall AUC score for each dimension led us to determine a cut-off of 5 for both Actionability and Knowledge scores. This meant that scores of 5 or higher were assigned a value of 1, while those below this threshold received a value of 0.

   Following the scoring procedure, the text extracted from webpages underwent pre-processing. This involved the removal of 'stop words' and the tokenization of the remaining words, a common approach when pre-processing textual data for analysis (Kelly & Sharot, 2023; Kelley & Gillan, 2022). The pre-processed text served as the input variable for our model, while the binary Actionability and Knowledge ratings were used as the target variables.

The next stage involved transforming the input data into a format suitable for machine learning. We achieved this by applying the TfidfVectorizer to the input variable, converting the textual data into a numerical matrix of Term Frequency-Inverse Document Frequency (TF-IDF) features.

Subsequently, to ensure the independence of samples, we first separated the data based on unique individuals, making certain that ratings from a specific individual either fell into the training set or the test set, but never both. This initial separation based on unique individuals ensured that there was no overlap or data leakage between the training and testing datasets at the individual level, thereby preventing individual-specific patterns or biases from influencing the model's performance.

Once we ensured the independence of data at the individual level, we then moved forward to address the potential class imbalance in our dataset. We applied the RandomOverSampler with the 'minority' sampling strategy. Following this resampling, the separated data was further divided into training and testing sets, maintaining a proportionate representation of the target variables. This split was stratified according to the target variables, allocating 30% of the total data to the test set.

Model training was then executed using the Light GBM Python package. Three logistic regression models, dedicated to Actionability and Knowledge respectively, were trained using the designated training set, from which we extracted the models' feature coefficients.

Lastly, the performance of the models was evaluated using the test set. For this, we utilised the eval function in the Light GBM package, providing us a robust measure of how effectively our models could predict Actionability and Knowledge ratings in a practical context.

Storage & Presentation of Values: The computed Emotion, Actionability, and Knowledge scores were stored in a system database. These scores can be subsequently distributed to users, system tools (e.g., browser plugins), or third parties (e.g., search engines). To maintain up-to-date scores, the process can be repeated periodically whenever the webpage content changes, its formatting is altered, or on a recurring interval basis (e.g., daily).

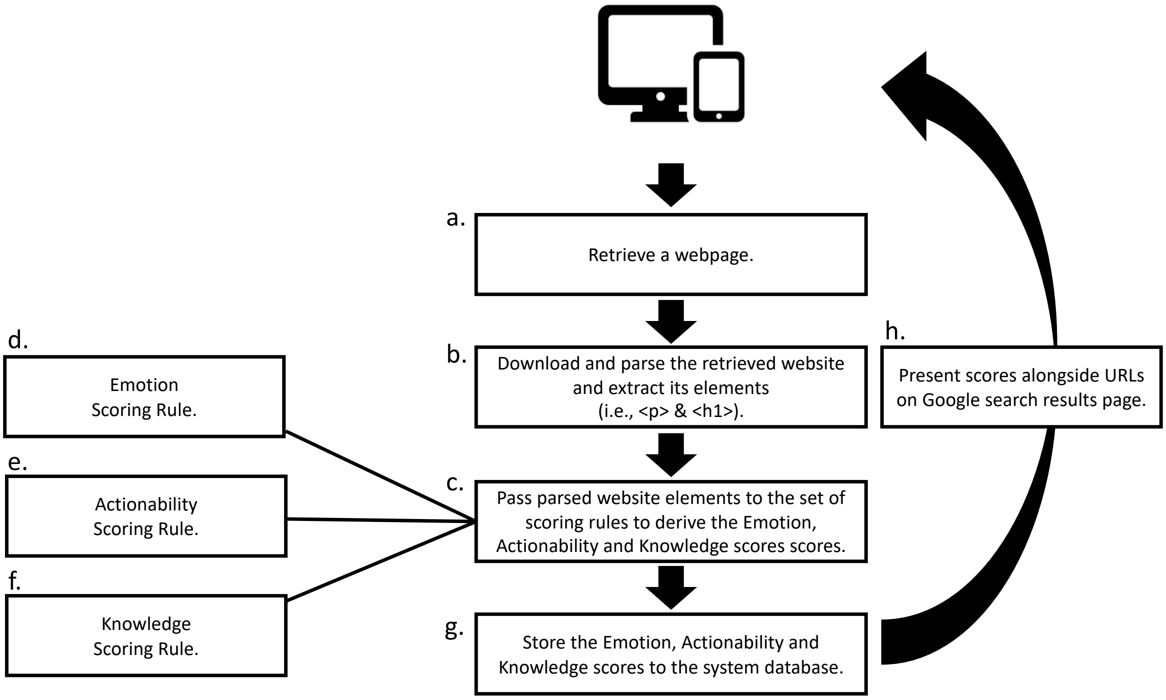

Figure 5.2. Process of Tool. The figure shows a visual representation of the tool's process. **(a)** the URLs of the webpages that users are exposed to on a Google search results page are retrieved. **(b)** Then, the HTML header and paragraph text of the webpages was downloaded and prepared for analysis. **(c-e)** The scoring rules for Emotion, Actionability, and Knowledge are defined and applied to the text. **(g)** The computed scores are stored in a database for subsequent distribution to users via a plugin. **(h)** The plugin presents the scores to users in real time enabling users to make informed decisions and adjust their information-consumption tendencies.

Model Performance: Whether information will help guide a person's action, increase their understanding of a topic or is perceived as positive will obviously alter from person to person. However, it is possible that ***on average*** some webpages contain information that is more likely to guide a person's action and/or increase their understanding on the topic. To test whether there is good reliability of these measures we did the following:

Emotion Model Performance Metrics

To assess the reliability of the VADER sentiment analysis tool in quantifying webpage sentiments, we conducted an experiment involving 500 human participants, recruited via Prolific's online recruitment platform (www.prolific.co). Each of these participants was tasked with freely browsing the internet and rate the sentiment of five webpages. Specifically, they were asked to rate *"how positive the information is on the webpage"* and *"how negative the information is on the webpage"*, on a scale from 1 ('not at all') to 6 ('very much'). Their ratings were then converted into an overall sentiment score: we calculated this score by subtracting the negative sentiment rating from the positive rating.

When we compared these human-generated sentiment scores with the VADER compound scores, we found a significant agreement (ICC = 0.712, $p < 0.001$). This strongly suggests that VADER's lexicon-based scoring method closely aligns with human subjective evaluations of webpage sentiment, effectively mirroring the emotional responses that people have when assessing the content of these webpages.

Actionability & Knowledge Model Performance Metrics

The performance of the logistic regression models were evaluated using several metrics. The precision, accuracy, and F1 score for each class were calculated, providing a comprehensive understanding of the model's ability to accurately classify the data. The results are shown in **Table 5.1**.

Table 5.1. *Performance Metrics of the Logistic Regression Models Predicting Actionability and Knowledge labels (i.e., 0 = low Actionability/Knowledge; 1 = High Actionability/Knowledge) of webpages.*

Model	Class	Precision	Recall	F1-Score	Support
Actionability	0	0.65	0.61	0.63	296
	1	0.72	0.74	0.73	389
	Accuracy	-	-	0.69	685
Knowledge	0	0.51	0.38	0.43	233
	1	0.72	0.81	0.76	452
	Accuracy	-	-	0.66	685

Precision is a measure of the proportion of correct positive predictions out of all positive predictions made by the model. In other words, it tells us how many of the instances predicted as positive (i.e., 1 = High Actionability/Knowledge) by the model are actually true positives.

$$Precision = True\ Positives \ / \ (True\ Positives + False\ Positives) \qquad (2)$$

For the Actionability model, the precision for class 0 was 0.65 (65%), indicating that 65% of the instances predicted as class 0 (i.e., low Actionability) were actually class 0. The precision for class 1 was 0.72 (72%), meaning that 72% of the instances predicted as class 1 were actually class 1, while for the Knowledge model, the precision for class 0 (i.e., low Knowledge) was 0.51 (51%), and for class 1, it was 0.72 (72%).

Recall (sensitivity) is a measure of how many of the actual positive instances (class 1) were correctly predicted by the model. It tells us the proportion of true positives that the model identified.

$$Recall = \frac{True\ Positives}{True\ Positives + False\ Negatives} \tag{3}$$

For the Actionability model, the recall for class 0 was 0.61 (61%), indicating that 61% of the actual class 0 instances were correctly predicted as class 0. The recall for class 1 was 0.74 (74%), meaning that 74% of the actual class 1 instances were correctly predicted as class 1, while for the Knowledge model, the recall for class 0 was 0.38 (38%), and for class 1, it was 0.81 (81%).

The F1-score is a single metric that balances both precision and recall, providing a more comprehensive evaluation of the model's performance. It is useful when you want to find a balance between precision and recall, especially in cases Of uneven class distribution.

$$F1\ score = 2\ \frac{Precision\ \times\ Recall}{Precision\ +\ Recall} \tag{4}$$

For the Actionability model, the F1-score for class 0 was 0.63, and for class 1, it was 0.73, while for the Knowledge model, the F1-score for class 0 was 0.43, and for class 1, it was 0.76.

Accuracy is the overall performance metric that measures how many instances were correctly classified by the model out of the total number of instances.

$$Accuracy = \frac{True\ Positives + True\ Negatives}{Total\ Instances} \tag{5}$$

For the Actionability model, the accuracy score was 0.69, indicating that the model correctly classified 69% of the instances in the test dataset, while for the Knowledge model, the accuracy score was 0.66 (or 66%).

ROC Curve Analysis and AUC Value

To further assess the performance of the logistic regression model, a Receiver Operating Characteristic (ROC) curve analysis was conducted. The ROC curve is a graphical representation of the true positive rate (sensitivity) against the false positive rate (1-specificity) at various decision threshold levels. The Area Under the Curve (AUC) value, which represents the overall performance of the classifier, was also calculated.

The AUC value obtained for the Actionability logistic regression model was 0.74, and for the Knowledge logistic regression model was 0.65. For reference, an AUC value of 0.5 represents a random classifier, whereas a value of 1 indicates a perfect classifier. Thus, the AUC values indicate that the Actionability model demonstrated good discriminative ability, while the Knowledge model showed average capability in distinguishing between their respective high and low levels of based on the webpage text data.

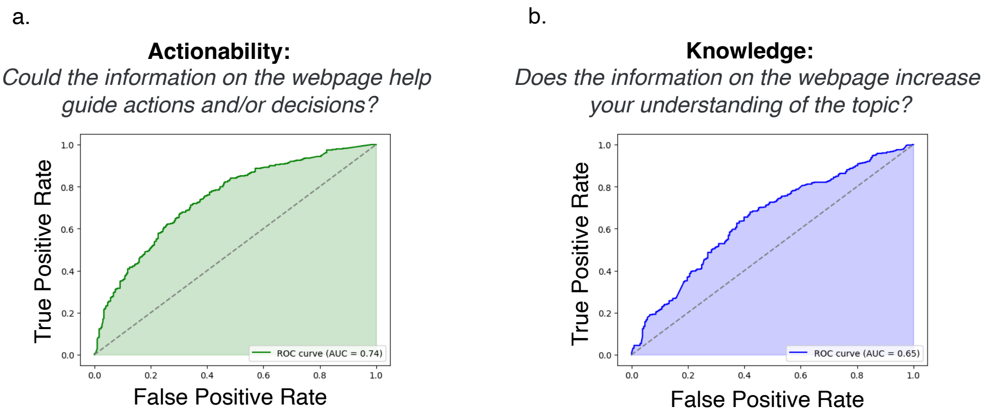

Figure 5.3. Performance Metrics of Actionability and Knowledge Models. We trained two logistic regression models to predict the (**a**) *Actionability* and (**b**) *Knowledge* classification of information on webpages. To do so, we asked participants to freely browse the internet and rate 5 webpages each on the following ratings: (**a**) *'Could the information on the webpage help guide actions and/or decisions?'*, and (**b**) *'Does the information on the webpage increase your understanding of the topic?'* on a 6-point scale, with 1 representing the lowest level of Actionability or Knowledge, and 6 representing the highest level. A binary threshold of 5 was applied to both Actionability and Knowledge scores: assigning a value of 1 if the value was equal to or greater than the threshold, and 0 otherwise. **(a&b)** Here we computed a Receiver Operating Characteristic (ROC) curve and Area Under the Curve (AUC) for the three logistic regression models. The ROC curve demonstrates the true positive rate (sensitivity) against the false positive rate (1-specificity) at various decision threshold levels. The AUC value for (**a**) Actionability is 0.74, and for (**b**) Knowledge is 0.65, which indicate that the Actionability model demonstrated good discriminative ability, while the Knowledge model showed average capability in distinguishing between their respective high and low levels of based on the webpage text data.

Power Analysis: The primary aim of this analysis was to determine the necessary number of participants needed to achieve our target AUC value of 0.8 (i.e., very good discrimination between classes) for both Actionability and Knowledge, assuming each participant rates 5 websites. To inform our estimates, we employed Kernel Density Estimation (KDE) to mirror the distribution of Actionability and Knowledge ratings based on their observed values. Our KDE-informed simulation suggests that reaching our desired AUC of 0.8 would necessitate collecting data from roughly 1317 individuals for Actionability ratings and 3418 individuals for Knowledge ratings. These estimates will guide our upcoming data acquisition and sampling endeavours.

5.3 Discussion

In the digital era, search engines have become an integral part of information gathering, with billions of queries submitted daily. Despite their prevalence, traditional search algorithms often fail to align with users web-browsing goals. To bridge this gap, we developed a unique tool, akin to nutritional labels for web content, to empower users to make informed decisions about the information they consume online.

Our tool, realised as a Google Chrome plugin, applies natural language processing (NLP) techniques to assign ‘Actionability’, ‘Knowledge’, and ‘Emotion’ scores to webpages. We based these criteria on empirical research suggesting these properties are pivotal to users’ information-seeking motives (Sharot & Sunstein, 2020; Kelly & Sharot, 2021; Cogliati Dezza et al., 2022). These scores are prominently displayed in Google search results (**see Figure 5.1**), guiding users in their web-browsing journey to align the information consumed more closely with their individual goals and preferences.

In terms of performance, the outcome metrics of the models (e.g., accuracy and AUC) were good for Actionability and average for Knowledge (Madrekar, 2010). This demonstrates the tool's effectiveness in classifying webpages according to their 'Actionability' and 'Knowledge' properties. Based on a power analysis, we have deduced that an increased sample size of approximately 1317 individuals for ‘Actionability’ ratings and 3418 for ‘Knowledge’ ratings should improve our models' performance, nudging them closer to our target AUC of 0.8. This insight will shape our subsequent data collection efforts. For Emotion, we observed good agreement between the algorithm's scoring of webpages and users' actual ratings. Together, these results suggest that despite the subjectivity and individual differences in interpretation, our tool was able to capture a shared perception. The scores can thus offer valuable guidance for users' online information consumption, allowing them to engage with information that aligns with their goals and preferences.

In the pursuit of continuously enhancing user experience and the tool's functionality, we have several potential directions to explore:

1**. Sort by Function:** We propose to allow users to reorder search results according to their preferred metrics. For example, presenting links in order from the most actionable to the least actionable. Such a feature could add an extra layer of customisability and empower users to tailor their information exposure according to their needs or preferences.

2. **Filter by Function:** Building on the ‘Sort by Function’, we suggest incorporating a filtering mechanism that allows users to eliminate search results based on one or multiple scores. Users might, for example, wish to exclude links with Knowledge and Actionability scores equal to or below 20. This approach could also be adapted as a parental tool, helping to guide children's online exposure.

3. **Track Web-Browsing Patterns Over Time:** Similar to apps that track physical activity or calorie intake, we envisage a feature that allows users to monitor their web-browsing patterns in relation to the three scores over time. This could provide valuable insights into their information-consumption tendencies and offer them the opportunity to adjust their browsing habits accordingly.

These proposed enhancements aim to augment user control over their web browsing, which may promote healthier, more constructive engagement with online content.

In addition, our current plugin focuses solely on analysing text and does not assess images and videos. While the results obtained suggest text analysis effectively reflects users' webpage ratings, we aim to broaden the tool's capabilities and include a diverse range of media types, such as images, videos, and other multimedia formats. By embracing these varied forms of media analysis, our goal is to create a more comprehensive and powerful tool that can assess all webpages and offer users a richer web-browsing experience.

The next step is to make the tool available to a diverse group of subjects to test (i) whether people select to use the tool (i.e. does exposure to the scores lead to changes their web-browsing patterns) and (ii) whether using the tool improves people's mood, subjective sense of knowledge enhancement and sense of empowerment.

Finally, the development and deployment of our Google Chrome plugin, while offering potentially significant benefits in improving online information consumption, also present potential harms that warrant careful consideration. A primary concern is the risk of misinformation or over-reliance on the tool, where users might accept the provided scores without critical evaluation, potentially leading to decisions based on inaccurate or biased information. The inherent subjectivity in scoring webpages for actionability, knowledge, and emotion could introduce biases, skewing content representation and possibly reinforcing existing biases or echo chambers. Additionally, the emphasis on sentiment analysis raises concerns about the impact on mental health, as continuous exposure to negatively valenced content could influence users' mental health. Finally, there's a risk that the tool, especially if expanded to include metrics like political sentiment, could inadvertently foster filter bubbles. This scenario, where users encounter only content that aligns with their existing viewpoints, might restrict access to a broad spectrum of opinions and information. Addressing these potential harms is crucial to ensure the tool's responsible and beneficial use in navigating the vast landscape of online information.

One way to address these concerns is establishing and adhering to accuracy benchmarks. This involves setting clear performance standards for the precision and recall of the tool's algorithms to minimise the risk of misinformation. Regularly reviewing and adjusting these benchmarks based on real-world usage data ensures that the tool remains reliable and effective. This process is essential not only for maintaining the integrity of the tool but also for building and retaining user trust, especially in an era where digital misinformation can have significant real-world consequences.

To conclude, our tool provides a novel solution to the shortcomings of traditional search algorithms by equipping users with an intuitive scoring system to assess web content. By integrating user-driven properties into search results, it enhances the browsing experience and facilitates more goal-oriented and effective online information consumption.

5.4 References

Caplin, A., & Leahy, J. (2001). Psychological expected utility theory and anticipatory feelings. The Quarterly Journal of Economics, 116(1), 55-79.

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, 115(31), E7255-E7264.

Dezza, I. C., Maher, C., & Sharot, T. (2022). People adaptively use information to improve their internal states and external outcomes. Cognition, 228, 105224.

Golman, R., Hagmann, D., & Loewenstein, G. (2017). Information avoidance. Journal of economic literature, 55(1), 96-135.

Kelley, S. W., & Gillan, C. M. (2022). Using language in social media posts to study the network dynamics of depression longitudinally. Nature communications, 13(1), 870.

Kelly, C. A., & Sharot, T. (2021). Individual differences in information-seeking. Nature communications, 12(1), 7062.

Kelly, C. A., & Sharot, T. (2023). Knowledge-Seeking Reflects and Shapes Well-Being.

Kemp. S. (2022). Digital 2022: Global Overview Report — DataReportal – Global Digital Insights. https://datareportal.com/reports/digital-2022-global-overview-report

Kőszegi, B. (2010). Utility from anticipation and personal equilibrium. Economic Theory, 44, 415-444.

Loewenstein, G. (1994). The psychology of curiosity: A review and reinterpretation. Psychological bulletin, 116(1), 75.

Mandrekar, J. N. (2010). Receiver operating characteristic curve in diagnostic test assessment. Journal of Thoracic Oncology, 5(9), 1315-1316.

Rainie, L., & Anderson, J. (2017). Code-dependent: Pros and cons of the algorithm age.

Richardson, L. (2007). Beautiful soup documentation. 2007. Retrieved on Mar, 24, 2020.

Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. Nature Human Behaviour, 4(1), 14-19.

Stigler, G. J. (1961). The Economics of Information. Journal of Political Economy. https://doi.org/10.1086/258464

Chapter 6

General Discussion

The studies comprising this thesis, which include Chapters 2-4, address three primary objectives. The first objective, discussed in Chapter 2, is to test an integrative theory of information-seeking proposed by Sharot and Sunstein (2020). The second objective, explored in Chapters 2-4, focuses on investigating the relationship between information-seeking patterns and well-being. Specifically, these studies examine whether information-seeking patterns in various contexts (i.e., information about personal-traits, web browsing, under stress) are linked to well-being. I hypothesised that observed relationships between information-seeking patterns and well-being would be bi-directional, and this is tested in Chapter 3. Chapter 5 introduces a third objective of the thesis, to develop an empirically informed tool (based on the findings from Chapters 2-4) to facilitate people's information-seeking online. In this discussion, I will present the strengths, limitations, applications, and future directions of this work.

An Integrative theory of Information-Seeking

In Chapter 2, I observed that an integrative theory of information-seeking proposed by Sharot and Sunstein (2020), which combines the three key empirical factors that drive information-seeking behaviour best described in the literature:

Instrumental Utility: People's information-seeking decisions are motivated by how useful the information will be in guiding their actions.
Hedonic Utility: People's information-seeking decisions are driven by the emotional impact that the information is expected to have on them.
Cognitive Utility: People seek information that is relevant to topics or concepts they think about often.

Importantly, I also found that a model which incorporates these three motives best explained individuals' information-seeking choices. This model was validated across four separate studies and three distinct domains: personal traits, finance, and health, suggesting its general applicability.

While the findings above highlight important motives for information-seeking behaviour, it's important to acknowledge that this might not encompass all factors. People may have various motives for seeking information, such as the desire to consume exciting content, which our study does not directly assess. Indeed, many algorithms are currently being developed to identify individual information preferences, aiming to curate content that aligns with these specific interests. However, it's doubtful that such a nuanced focus on individual motives would effectively explain information-seeking motives at a population level due to the potentially diverse nature of these motives among individuals. Therefore, the approach employed in this thesis, assessing the broad motives, provides an overarching understanding of the primary reasons behind information-seeking behaviour.

Building on this understanding, Chapter 4 investigated how context and environment further shaped such motives. My observations indicate that during global and personal stressful events, there's a significant shift towards seeking more instrumental information, as evidenced by the rise in *'How'* searches during these periods. Given this insight, it may explain the fact that uncertainty – a factor often linked to information-seeking in numerous studies (Afifi & Weiner, 2004; Berlyne, 1957; Kreps & Porteus, 1978; Nickerson, 1998; Klayman & Ha, 1987; Kappes et al., 2020; Stigler, 1961; Kobayashi et al., 2019; Charpentier et al., 2018; Hirshleifer & Riley, 1979; van Lieshout et al., 2018; Trudel et al., 2021) – was not identified as a motive in the winning models of Chapter 2. In other words, context and environment may change what motive is dominant.

Relationship between Information-Seeking and Mental Health

In recent years, as people spend more time online, the need to investigate the relationship between online activity and mental health has become imperative (Kemp, 2021). In this thesis, I've shown that the relationship between information-seeking behaviour and mental health is bi-directional and causal. Moreover, I presented that relationship between information-seeking and mental health can be context dependent. In particular, stressful events can trigger different information-seeking patterns, which in turn, may have implications on well-being.

In Chapter 3, I examined the relationship between individuals' web browsing patterns and their well-being, a matter that has grown increasingly significant in the 'information era'. My analysis revealed a positive correlation between the negativity of the information consumed by people and the intensity of mental health symptoms they self-reported through questionnaires. This effect was also bi-directional - participants who reported worse mood prior to browsing tended to consume more negative content online and exposure to negative content was in turn associated with worse mood (controlling for pre-browsing mood). Importantly, I established the causality of this relationship by first exposing participants to either negative or neutral webpages, where I found that exposure to negative webpages resulted in worse mood, and this change in mood then led to more browsing of negatively valence information. Together, these findings suggest a feedback loop; low mood leads to the consumption of more negative information online which in turn leads to worse mood and so on. By integrating information-seeking and sharing methodologies, we can achieve a broader account of the relationship between online well-being and online behaviour.

This work is innovative in its approach of examining the link between the information browsed and mental health. Previous research in this area has focused on analysing specific search engine queries rather than the actual text on webpages visited (Ayers et al., 2021; Gunnell et al., 2015). This traditional approach monitors certain keywords, such as "therapist" or "Prozac", to infer changes in population mental health. This method may be limited in its ability to assess an individual's mental health, as it only provides a limited dataset based on a few keywords, which would be used by individuals who are already aware of their symptoms and seek help.

The rationale of our approach, which is to quantify the affective properties of text, aligns with studies that demonstrate a relationship between the affective properties of shared content and mental health, such as social media posts (De Choudhury et al., 2017; Chancellor et al., 2019; Kelley & Gillan, 2022; Eichstaedt et al., 2018). This suggests a potentially intriguing overlap between the mechanisms governing information-seeking and those governing information-sharing (e.g., Vellani et al., 2022). In other words, is there a correlation between the information that people share online and the information they consume, and vice versa? Indeed, the size of the effects reported here are comparable to those found for the relationship between mental health and information-sharing (Kelley & Gillan, 2022) (as well as for those found between mental health and frequency of social media use; Twenge et al., 2017; Yoon et al., 2019; McCrae et al., 2017; Vahedi & Zannella, 2021). An advantage of the current approach is that it does not necessitate that people share information online, an activity that is clearly prevalent, but less so than online information-seeking. However, combining the two approaches would offer a more holistic account of the relationship between online behaviour and mental health.

Building on this understanding, it is vital to consider the role of algorithms in shaping the information we encounter, particularly regarding its effects on our mental health. Digital media platforms, including social media, utilise attention-grabbing elements driven by artificial intelligence algorithms to promote user engagement and interaction (Billieux, 2015; Christakis, 2019). However, there is a notable gap in our understanding of how these algorithms impact mental health. Moreover, it is important to investigate how algorithms on such platforms impact our online behaviour on lighter algorithmic platforms such a interacting with search browsers (e.g., Google and Bing). However, in this thesis, I did not assess behaviour on password-protected sites such as social media. Given the significant time individuals spend on these platforms (Kemp, 2021) and their algorithm-driven nature, investigating the potential for self-reinforcing feedback loops is critical. It is possible that algorithms tailored to users' past behaviours could perpetuate affective states by promoting specific types of content, such as those with negative valence. This might amplify the feedback loops we've noted, significantly influencing user experiences across the digital landscape. In contrast, algorithms may also be a source of noise, in the sense that they alter peoples' natural search intentions. If that is the case, the relationship between mental health and self-driven searches may even be greater than we observe here.

To assess the entire media space of online behaviour, I could utilise the *Screenomics* methodology (Reeves et al., 2021), a valuable approach for enhancing our understanding of online behaviour. This dataset provides a granular view of digital engagement by recording screen interactions every five seconds over the course of a year, activated whenever a user's device is in use (~70 TB of image time series data). While providing screen data, participants also completed bi-weekly self-reports about their mental health, including depression, anxiety, stress, and sleep. The temporal density and multimodal nature of the data enables a holistic analysis of how users interact with online content in everyday life, thereby improving ecological validity.

Working with these data, I would be able to obtain a holistic view of the digital environment's effects on users, and how algorithms on different platforms differentially influence browsing patterns and impact mental health.

Intervention to Facilitate Online Information-Seeking and Improve Well-Being

In Chapter 3, building upon our findings regarding the bi-directional relationship between negative information consumption online and mood, I investigated how informing individuals about the emotional impact of webpages could influence their web-browsing choices. This approach, grounded in the hypothesis that awareness of a webpage's likely affective impact could direct individuals towards less negatively-valenced content, was validated by my observations. These showed that pre-emptive cues about a webpage's emotional impact effectively steered browsing patterns away from negative content. Such shifts in browsing behaviour hold the potential to incrementally improve individuals' mental health over time, though further exploration is needed to confirm this. I propose that this approach, in conjunction with existing intervention tools such as screen time awareness applications (Kovacs et al., 2021) and digital phenotyping methods (Reece & Danforth, 2017; Guntuku et al., 2020; Valdez et al., 2020; De Choudhury et al., 2013; Kelley & Gillan, 2022; Eichstaedt et al., 2018; Torous et al., 2016), could significantly help mitigate the negative impact of internet use on mental health.

Moreover, in many cases it would be obviously suboptimal to solely base information-consumption decisions on the affective properties of information. For example, if someone searches for information on whether smoking causes cancer, the most positive link may not necessarily be the wise choice. Thus, as mentioned, we do not envision the intervention described here as a stand-alone tool. Rather, by providing users with affective labels in addition to other labels, such as the reliability and instrumental utility of information, users can make more informed decisions that align with their current goals. For instance, users may want to prioritise the instrumental utility of information in one situation and prioritize their mood in another by focusing on affective labels.

Indeed, in Chapter 5, building on insights from Chapters 2-4, I presented a tool that's designed to facilitate online web-browsing. Currently, search results are shaped by algorithms that don't always match what users are looking for. This can lead to people spending a lot of time on unhelpful information, which can sometimes impact their well-being negatively. Our tool, which works as a Google Chrome plugin, is analogues to a 'nutrition label' for web content. Instead of showing calories and fat, our labels provide users a heads-up about what kind of information a webpage contains before they click on it.

Specifically, the tool quantifies three key properties of a webpage: (i) actionability, which measures how useful the text is for taking actions (i.e., Instrumental Utility); (ii) knowledge enhancement, which reflects how much a text can help users understand a topic (i.e., Cognitive Utility); and (iii) emotion, which quantifies the text's positive/negative tone (i.e., Hedonic Utility). These properties are derived from insights from Chapters 2-4, highlighting the motives for information-seeking.

To verify the effectiveness of the tool, further studies are necessary. My focus will be on evaluating how it influences users' online behaviour and whether this leads to enhanced browsing experiences in line with their web-browsing goals and well-being. To that end, participants will be invited to install our plugin and use it for a specified duration. Additionally, I'll measure various user experience metrics to refine the tool's usability. Feedback on preferred information features will be sought from participants to better cater to their browsing requirements (e.g., excitement of information etc.). This approach aims to uncover the personal motivations of users, providing them with a tailored online experience.

Applications

The research presented in this thesis, encompassing Chapters 2-5, has several practical applications. These applications can be broadly categorised into a few key areas:

1. Enhancing Information Integration: Recognising individual differences in information-seeking behaviour, particularly in how people value the three motives (Instrumental, Hedonic, and Cognitive Utilities), can lead to more effective information consumption. For instance, policymakers can enhance the effectiveness of critical information dissemination, such as voting procedures, by tailoring their communications to encompass these utilities, either in a single message or a series of messages. This approach could significantly boost the overall impact of their campaigns.
2. Guiding Algorithm Design and Content Curation: The insights gained from this research can be beneficial in shaping the algorithms used by search engines and social media platforms. Understanding the motives behind information-seeking and the effects of content on well-being can lead to the creation of algorithms that not only cater to user preferences but also promote content that is beneficial to mental health.
3. Advancing Awareness and Education on Digital Well-being: The methodologies employed in this research, particularly those analysing web-browsing patterns and their emotional impact, are valuable tools for education and awareness. They can be effectively utilised in educational settings as well as in the general population, enabling both children and adults to deepen their understanding of how online information impacts our well-being. This approach will help to better educate about the digital landscape and its effects on mental health.
4. Furthering Academic Research: Finally, this thesis lays the groundwork for further academic research in various fields, including psychology, information science, and computer science. It opens new avenues for exploring the complex interplay between online behaviour, information-seeking, and well-being, thereby enriching the academic discourse in these areas.

In essence, the applications of this research are broad and multifaceted, impacting everything from individual browsing habits to large-scale digital policymaking. It contributes significantly to our understanding of the digital world and its complex relationship with human behaviour and mental health.

Concluding Remarks

The research undertaken in this thesis, spread across Chapters 2-5, provides a comprehensive understanding of the intricate dynamics inherent to information-seeking behaviours and their diverse relationship to well-being. This work holds relevance for the academic community and policy makers, offering valuable insights into the complexities of information-seeking behaviour. Furthermore, it extends its significance to the broader public as it introduces an intervention aimed at facilitating informed information-seeking decisions.

Central to the narrative is the understanding that our information-seeking choices aren't passive or isolated acts. Instead, they reflect our psychological needs and desires, and shape our states of mind. Chapter 2's examination of Sharot and Sunstein's (2020) integrative theory serves as a testament to this, offering a robust framework that highlights the motives for why we seek information: to gain Instrumental, Hedonic, and Cognitive Utility. The variation in individual preferences in relation to these utilities highlights the notion that our information-seeking choices are an extension of our traits and states, given their relative stability across domains and over time.

Chapter 3 supports this narrative further by revealing the bi-directional relationship between online web-browsing patterns and well-being. Specifically, the results suggest a feedback loop of mood and content consumption, emphasising the impact of our digital engagements, especially in an era where much of our time is spent online. The complexities of this relationship are further nuanced under different contexts, for example, under stress as emphasised in Chapter 4. The tendency to seek action-oriented queries during stressful periods can be seen both as an adaptive mechanism and a reflection of the human instinct to seek clarity amidst chaos. While, Chapter 3 underscores the idea that individuals, when gravitating towards information-seeking about frequently contemplated self-concepts, might enhance their mental well-being by achieving cognitive closure.

Informed by these insights is the development and introduction of the web-browsing tool in Chapter 5. This tool, albeit in its early stages, shows a lot of promise, illuminating the potential for a paradigm shift in how we navigate the online environment. By providing users with scores of webpages' Actionability, Knowledge, and Emotional characteristics, it seeks to elevate the user experience from mere browsing to informed, purpose-driven engagement.

In conclusion, this thesis contributes to the scientific understanding of information-seeking and its multidimensional impact on well-being. It raises the conceptual and methodological bar for future research in this critical area and offers practical applications that could shape the way we interact with the digital world. As we continue to navigate the information age, where our online and offline realities are increasingly blurred, this research stands as a cornerstone, urging us to consider not just what information we seek, but also why we seek it, and how that, in turn, shapes our psychological well-being.

References

Afifi, W. A., & Weiner, J. L. (2004). Toward a theory of motivated information management. Communication Theory, 14(2), 167-190.

Ayers, J. W., Poliak, A., Johnson, D. C., Leas, E. C., Dredze, M., Caputi, T., & Nobles, A. L. (2021). Suicide-Related Internet Searches During the Early Stages of the COVID-19 Pandemic in the US. JAMA Network Open, 4(1). https://doi.org/10.1001/JAMANETWORKOPEN.2020.34261

Berlyne, D. E. (1957). Uncertainty and conflict: a point of contact between information-theory and behavior-theory concepts. Psychological Review, 64(6p1), 329.

Billieux, J., Maurage, P., Lopez-Fernandez, O., Kuss, D. J., & Griffiths, M. D. (2015). Can disordered mobile phone use be considered a behavioral addiction? An update on current evidence and a comprehensive model for future research. Current Addiction Reports, 2(2), 156-162.

Chancellor, S., Birnbaum, M., Caine, E., Silenzio, V. & De Choudhury, M. A taxonomy of ethical tensions in inferring mental health states from Social media. in Proc. Conference on Fairness, Accountability, and Transparency (FAT*) (ACM, 2019).

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, 115(31), E7255-E7264.

Christakis, D. A. (2019). The challenges of defining and studying “digital addiction” in children. Jama, 321(23), 2277-2278.

De Choudhury, M., Gamon, M., Counts, S. & Horvitz, E. (2013). Predicting depression via social media. ICWSM 2, 128–137.

De Choudhury, M., Gamon, M., Counts, S. & Horvitz, E. (2013). Predicting depression via social media. ICWSM 2, 128–137.

Eichstaedt, J. C. et al. Facebook language predicts depression in medical records. Proc. Natl Acad. Sci. USA 115, 11203–11208 (2018).

Gunnell, D., Derges, J., Chang, S. Sen, & Biddle, L. (2015). Searching for suicide methods. Crisis, 36(5), 325–331. https://doi.org/10.1027/0227-5910/A000326

Guntuku, S. C., Sherman, G., Stokes, D. C., Agarwal, A. K., Seltzer, E., Merchant, R. M., & Ungar, L. H. (2020). Tracking mental health and symptom mentions on Twitter during COVID-19. Journal of general internal medicine, 35, 2798-2800.

Hirshleifer, J., & Riley, J. G. (1979). The Analytics of Uncertainty and Information- An Expository Survey. Source Journal of Economic Literature Journal of Economic Literature.

Kappes, A., Harvey, A. H., Lohrenz, T., Montague, P. R., & Sharot, T. (2020). Confirmation bias in the utilization of others’ opinion strength. Nature neuroscience, 23(1), 130-137.

Kelley, S. W., & Gillan, C. M. (2022). Using language in social media posts to study the network dynamics of depression longitudinally. Nature communications, 13(1), 870.

Kemp. S. (2021). Digital 2021: Global Overview Report — DataReportal – Global Digital Insights. https://datareportal.com/reports/digital-2021-global-overview-report

Klayman, J., & Ha, Y. W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. Psychological review, 94(2), 211.
Kobayashi, K., Ravaioli, S., Baranès, A., Woodford, M., & Gottlieb, J. (2019). Diverse motives for human curiosity. Nature human behaviour, 3(6), 587-595.
Kovacs, B., Caplan, N., Grob, S., & King, M. (2021). Social networks and loneliness during the COVID-19 pandemic. Socius, 7, 2378023120985254.
Kreps, D. M., & Porteus, E. L. (1978). Temporal resolution of uncertainty and dynamic choice theory. Econometrica: journal of the Econometric Society, 185-200.
McCrae, N., Gettings, S., & Purssell, E. (2017). Social media and depressive symptoms in childhood and adolescence: A systematic review. Adolescent Research Review, 2, 315-330.
Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. Review of general psychology, 2(2), 175-220.
Reece, A. G., & Danforth, C. M. (2017). Instagram photos reveal predictive markers of depression. EPJ Data Science, 6(1), 15.
Reeves, B., Ram, N., Robinson, T. N., Cummings, J. J., Giles, C. L., Pan, J., ... & Yeykelis, L. (2021). Screenomics: A framework to capture and analyze personal life experiences and the ways that technology shapes them. Human–Computer Interaction, 36(2), 150-201.
Sharot, T., & Sunstein, C. R. (2020). How people decide what they want to know. Nature Human Behaviour, 4(1), 14-19.
Stigler, G. J. (1961). The Economics of Information. Journal of Political Economy. https://doi.org/10.1086/258464
Torous, J., Kiang, M. V., Lorme, J., & Onnela, J. P. (2016). New tools for new research in psychiatry: a scalable and customizable platform to empower data driven smartphone research. JMIR mental health, 3(2), e5165.
Trudel, N., Scholl, J., Klein-Flügge, M. C., Fouragnan, E., Tankelevitch, L., Wittmann, M. K., & Rushworth, M. F. (2021). Polarity of uncertainty representation during exploration and exploitation in ventromedial prefrontal cortex. Nature Human Behaviour, 5(1), 83-98.
Twenge, J. M. (2017). Have smartphones destroyed a generation. The Atlantic, 9, 2017.
Vahedi, Z., & Zannella, L. (2021). The association between self-reported depressive symptoms and the use of social networking sites (SNS): A meta-analysis. Current Psychology, 40, 2174-2189.
Valdez, D., Ten Thij, M., Bathina, K., Rutter, L. A., & Bollen, J. (2020). Social media insights into US mental health during the COVID-19 pandemic: Longitudinal analysis of Twitter data. Journal of medical Internet research, 22(12), e21418.
van Lieshout, L. L., Vandenbroucke, A. R., Müller, N. C., Cools, R., & de Lange, F. P. (2018). Induction and relief of curiosity elicit parietal and frontal activity. Journal of Neuroscience, 38(10), 2579-2588.
Vellani, V., & Sharot, T. (2022). How people decide when to inform others.
Yoon, S., Kleinman, M., Mertz, J., & Brannick, M. (2019). Is social network site usage related to depression? A meta-analysis of Facebook–depression relations. Journal of affective disorders, 248, 65-72.

Supplementary Materials

Chapter 2

Supplementary Table 2.1. *Study 1: Average scores of ratings and their correlations.*

Mean of task variables (standard deviation).
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test **comparing mean to mid-point of their scale**).

	Information-Seeking Choice ***(-3 to +3)***	Usefulness ***(-3 to +3)***	Thought Frequency ***(-3 to +3)***	Feelings to Know ***(-3 to +3)***	Feelings Never to Know ***(-3 to +3)***	Feelings to Know - Feelings Never to Know ***(-6 to +6)***	Expected rating of others ***(Negative traits reverse scored) -3 to +3)***	Confidence in estimation ***(-3 to +3)***
Task Ratings Mean (SD)	0.43** (1.30)	0.60*** (1.19)	0.15 (0.85)	0.73*** (0.89)	0.06 (0.92)	0.67*** (1.35)	0.95*** (1.07)	1.55*** (0.78)

Mean Pearson R between variables calculated for each participant and then averaged across participants (standard deviation)**.**
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test comparing the mean R across subjects to zero).

	Information-Seeking Choice	Usefulness	Thought Frequency	Feelings to Know	Feelings Never to Know	Feelings to Know - Feelings Never to Know	Expected rating	Confidence in estimation
Information-Seeking Choice Mean (SD)	-							
Usefulness Mean (SD)	0.21*** (0.23)	-						
Thought Frequency Mean (SD)	0.21*** (0.24)	0.21*** (0.23)	-					
Feelings to Know Mean (SD)	0.23*** (0.23)	0.31*** (0.29)	0.16*** (0.27)	-				
Feelings Never to Know Mean (SD)	-0.13*** (0.23)	-0.24*** (0.31)	-0.10** (0.25)	-0.25*** (0.41)	-			
Feelings to Know - Feelings Never to Know Mean (SD)	0.23*** (0.24)	0.34*** (0.29)	0.17*** (0.26)	0.82*** (0.20)	-0.70*** (0.21)	-		
Expected rating of others Mean (SD)	0.19*** (0.26)	0.15*** (0.31)	0.07** (0.22)	0.40*** (0.32)	-0.21*** (0.28)	0.38*** (0.31)	-	

Confidence in estimation Mean (SD)	0.05 (0.22)	0.02 (0.38)	0.10*** (0.20)	0.25*** (0.29)	-0.02 (0.29)	0.18*** (0.29)	0.29*** (0.31)	-

Information-Seeking Choice Rating: -3 ('*definitely don't want to know*') to +3 ('*definitely want to know*'); **Usefulness:** -3 ('*not useful*') to +3 (*very useful*); **Thought Frequency:** -3 ('*never*') to +3 ('very *often*'); **Feelings to Know:** -3 ('*very bad')* to +3 ('*very good*'); **Feelings Not to Know:** -3 ('*very bad')* to (+3 '*very good*'); **Expected Rating of Others:** -3 (*'not at all this trait')* to +3 (*'very much this trait';* scores were reversed for negative valanced stimuli); **Confidence in Estimation:** -3 ('*not certain')* to +3 (*'very certain*').

Supplementary Table 2.2. *Study 2, Time 1: Average scores of ratings and their correlations.*

Mean of task variables (standard deviation).
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test **comparing mean to mid-point of their scale**).

	Information-Seeking Choice *(-3 to +3)*	Usefulness *(-3 to +3)*	Thought Frequency *(-3 to +3)*	Feelings to Know *(-3 to +3)*	Feelings Never to Know *(-3 to +3)*	Feelings to Know - Feelings Never to Know *(-6 to +6)*	Expected rating of others ***(Negative traits reverse scored)*** ***-3 to +3)***	Confidence in estimation *(-3 to +3)*
Task Ratings Mean (SD)	1.17*** (1.02)	0.75*** (1.01)	0.53*** (0.81)	0.67*** (0.72)	0.00 (0.86)	0.68*** (1.16)	0.93*** (0.66)	1.20*** (0.82)

Mean Pearson R between variables calculated for each participant and then averaged across participants (standard deviation)**.**
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test comparing the mean R across subjects to zero).

	Information-Seeking Choice	Usefulness	Thought Frequency	Feelings to Know	Feelings Never to Know	Feelings to Know - - Feelings Never to Know	Expected rating of others	Confidence in estimation
Information-Seeking Choice Mean (SD)	-							
Usefulness Mean (SD)	0.165*** (0.23)	-						
Thought Frequency Mean (SD)	0.12*** (0.21)	0.20*** (0.22)	-					
Feelings to Know Mean (SD)	0.18*** (0.22)	0.30*** (0.32)	0.09*** (0.24)	-				
Feelings Never to Know Mean (SD)	-0.12*** (0.22)	-0.23*** (0.31)	-0.08*** (0.24)	-0.20*** (0.40)	-			
Feelings to Know - Feelings Never to Know Mean (SD)	0.19*** (0.22)	0.32*** (0.30)	0.10*** (0.23)	0.83*** (0.16)	-0.68*** (0.24)	-		
Expected rating of others Mean (SD)	0.20*** (0.23)	0.15*** (0.26)	-0.02 (0.21)	0.35*** (0.29)	-0.14*** (0.27)	0.32*** (0.29)	-	
Confidence in estimation Mean (SD)	0.07*** (0.20)	0.02 (0.33)	0.11*** (0.21)	0.19*** (0.27)	0.01 (0.25)	0.13*** (0.27)	0.23*** (0.27)	-

Information-Seeking Choice Rating: -3 (‘*definitely don’t want to know*’) to +3 (‘*definitely want to know*’); **Usefulness:** -3 (‘*not useful*’) to +3 (*very useful*); **Thought Frequency:** -3 (‘*never*’) to +3 (‘very *often*’); **Feelings to Know:** -3 (‘*very bad’)* to +3 (‘*very good*’); **Feelings Not to Know:** -3 (‘*very bad’)* to (+3 ‘*very good*’); **Expected Rating of Others:** -3 (*'not at all this trait')* to +3 (*‘very much this trait';* scores were reversed for negative valanced stimuli); **Confidence in Estimation:** -3 (‘*not certain')* to +3 (*‘very certain*’).

Supplementary Table 2.3. *Study 2, Time 2: Average scores of ratings and their correlations.*

Mean of task variables (standard deviation).
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test **comparing mean to mid-point of their scale**).

	Information-Seeking Choice *(-3 to +3)*	Usefulness *(-3 to +3)*	Thought Frequency *(-3 to +3)*	Feelings to Know *(-3 to +3)*	Feelings Never to Know *(-3 to +3)*	Feelings to Know - Feelings Never to Know *(-6 to +6)*	Expected rating of others ***(Negative traits reverse scored) -3 to +3)***	Confidence in estimation *(-3 to +3)*
Task Ratings Mean (SD)	0.74*** (0.82)	0.56*** (0.93)	0.41*** (0.78)	0.52*** (0.65)	0.08 (0.75)	0.43*** (1.05)	0.72*** (0.67)	1.03*** (0.77)

Mean Pearson R between variables calculated for each participant and then averaged across participants (standard deviation)**.**
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test comparing the mean R across subjects to zero).

	Information-Seeking Choice	Usefulness	Thought Frequency	Feelings to Know	Feelings Never to Know	Feelings to Know - Feelings Never to Know	Expected rating of others	Confidence in estimation
Information-Seeking Choice Mean (SD)	-							
Usefulness Mean (SD)	0.18*** (0.20)	-						
Thought Frequency Mean (SD)	0.16*** (0.22)	0.21*** (0.23)	-					
Feelings to Know Mean (SD)	0.19*** (0.23)	0.30*** (0.31)	0.12*** (0.25)	-				
Feelings Never to Know Mean (SD)	-0.13*** (0.21)	-0.18*** (0.32)	-0.08*** (0.23)	-0.17*** (0.42)	-			
Feelings to Know - Feelings Never to Know Mean (SD)	0.20*** (0.21)	0.31*** (0.30)	0.12*** (0.24)	0.82** (0.17)	-0.67*** (0.23)	-		
Expected rating of others Mean (SD)	0.21*** (0.25)	0.14*** (0.28)	0.00 (0.23)	0.35*** (0.29)	-0.11*** (0.27)	0.31*** (0.29)	-	
Confidence in estimation Mean (SD)	0.02 (0.19)	0.01 (0.34)	0.11*** (0.21)	0.17*** (0.26)	0.04 (0.23)	0.11*** (0.24)	0.20** (0.27)	-

Information-Seeking Choice Rating: -3 ('*definitely don't want to know*') to +3 ('*definitely want to know*'); **Usefulness:** -3 ('*not useful*') to +3 (*very useful*); **Thought Frequency:** -3 ('*never*') to +3 ('very *often*'); **Feelings to Know:** -3 ('*very bad')* to +3 ('*very good*'); **Feelings Not to Know:** -3 ('*very bad')* to (+3 '*very good*'); **Expected Rating of Others:** -3 (*'not at all this trait')* to +3 (*'very much this trait';* scores were reversed for negative valanced stimuli); **Confidence in Estimation:** -3 ('*not certain ')* to +3 (*'very certain '*).

Supplementary Table 2.4. *Study 3, Time 1: Average scores of ratings and their correlations.*

Mean of task variables (standard deviation).
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test **comparing mean to mid-point of their scale**).

	Information-Seeking Choice *(-3 to +3)*	Usefulness *(-3 to +3)*	Thought Frequency *(-3 to +3)*	Feelings to Know *(-3 to +3)*	Feelings Never to Know *(-3 to +3)*	Feelings to Know - Feelings Never to Know *(-6 to +6)*	Confidence in estimation *(-3 to +3)*
Task Ratings Mean (SD)	1.12*** (0.90)	1.14*** (0.95)	0.67*** (1.04)	0.99*** (0.93)	-0.34*** (0.96)	1.33*** (1.39)	0.42*** (1.07)

Mean Pearson R between variables calculated for each participant and then averaged across participants (standard deviation)**.**
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test comparing the mean R across subjects to zero).

	Information-Seeking Choice	Usefulness	Thought Frequency	Feelings to Know	Feelings Never to Know	Feelings to Know - Feelings Never to Know	Confidence in estimation
Information-Seeking Choice Mean (SD)	-						
Usefulness Mean (SD)	0.37*** (0.29)	-					
Thought Frequency Mean (SD)	0.27*** (0.24)	0.30*** (0.25)	-				
Feelings to Know Mean (SD)	0.25*** (0.25)	0.37*** (0.36)	0.18*** (0.25)	-			
Feelings Never to Know Mean (SD)	-0.23*** (0.22)	-0.33*** (0.28)	-0.22*** (0.27)	-0.26*** (0.34)	-		
Feelings to Know - Feelings Never to Know Mean (SD)	0.29*** (0.24)	0.42*** (0.31)	0.23*** (0.24)	0.80*** (0.16)	-0.75*** (0.17)	-	
Confidence in estimation Mean (SD)	0.11*** (0.16)	0.17*** (0.26)	0.19*** (0.22)	0.10*** (0.24)	-0.05** (0.21)	0.09*** (0.20)	-

Information-Seeking Choice Rating: -3 ('*definitely don't want to know*') to +3 ('*definitely want to know*'); **Usefulness:** -3 ('*not useful*') to +3 (*very useful*); **Thought Frequency:** -3 ('*never*') to +3 ('very *often*'); **Feelings to Know:** -3 ('*very bad')* to +3 ('*very good*'); **Feelings Not to Know:** -3 ('*very bad')* to (+3 '*very good*'); **Confidence in Estimation:** -3 ('*not certain')* to +3 (*'very certain*').

Supplementary Table 2.5. *Study 3, Time 2: Average scores of ratings and their correlations.*

Mean of task variables (standard deviation).
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test **comparing mean to mid-point of their scale**).

	Information-Seeking Choice *(-3 to +3)*	Usefulness *(-3 to +3)*	Thought Frequency *(-3 to +3)*	Feelings to Know *(-3 to +3)*	Feelings Never to Know *(-3 to +3)*	Feelings to Know - Feelings Never to Know *(-6 to +6)*	Confidence in estimation *(-3 to +3)*
Task Ratings Mean (SD)	0.89*** (1.00)	0.79*** (1.12)	0.28* (1.12)	0.71*** (1.00)	-0.33** (0.95)	1.04*** (1.40)	0.09 (1.11)

Mean Pearson R between variables calculated for each participant and then averaged across participants (standard deviation).
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test comparing the mean R across subjects to zero).

	Information-Seeking Choice	Usefulness	Thought Frequency	Feelings to Know	Feelings Never to Know	Feelings to Know - Feelings Never to Know	Confidence in estimation
Information-Seeking Choice Mean (SD)	-						
Usefulness Mean (SD)	0.44*** (0.29)	-					
Thought Frequency Mean (SD)	0.37*** (0.23)	0.43*** (0.25)	-				
Feelings to Know Mean (SD)	0.28*** (0.24)	0.42*** (0.34)	0.26*** (0.29)	-			
Feelings Never to Know Mean (SD)	-0.30*** (0.25)	-0.32*** (0.40)	-0.28*** (0.30)	-0.29*** (0.42)	-		
Feelings to Know - Feelings Never to Know Mean (SD)	0.34*** (0.24)	0.42*** (0.33)	0.30*** (0.27)	0.82*** (0.17)	-0.75*** (0.19)	-	
Confidence in estimation Mean (SD)	0.14*** (0.17)	0.18*** (0.29)	0.25*** (0.19)	0.10** (0.30)	-0.07* (0.26)	0.09*** (0.25)	-

Information-Seeking Choice Rating: -3 ('*definitely don't want to know*') to +3 ('*definitely want to know*'); **Usefulness:** -3 ('*not useful*') to +3 (*very useful*); **Thought Frequency:** -3 ('*never*') to +3 ('very *often*'); **Feelings to Know:** -3 ('*very bad')* to +3 ('*very good*'); **Feelings Not to Know:** -3 ('*very bad')* to (+3 '*very good*'); **Confidence in Estimation:** -3 ('*not certain')* to +3 (*'very certain*').

Supplementary Table 2.6. *Study 4: Average scores of ratings and their correlations.*

Mean of task variables (standard deviation).
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test **comparing mean to mid-point of their scale**).

	Information-Seeking Choice *(-3 to +3)*	Usefulness *(-3 to +3)*	Thought Frequency *(-3 to +3)*	Feelings to Know *(-3 to +3)*	Feelings Never to Know *(-3 to +3)*	Feelings to Know - Feelings Never to Know *(-6 to +6)*	Expectations ***(Negative stimuli reverse scored) -3 to +3)***	Confidence in estimation *(-3 to +3)*
Task Ratings Mean (SD)	1.47*** (0.99)	1.13*** (0.88)	-0.15 (0.90)	0.89*** (0.84)	-0.29*** (0.77)	1.13*** (1.21)	0.40*** (0.69)	0.80*** (0.99)

Mean Pearson R between variables calculated for each participant and then averaged across participants (standard deviation)**.**
(***p<0.001, **p <0.01, *p<0.05 (two-sided): one-sample t-test comparing the mean R across subjects to zero).

	Information-Seeking Choice	Usefulness	Thought Frequency	Feelings to Know	Feelings Never to Know	Feelings to Know - Feelings Never to Know	Expected rating	Confidence in estimation
Information-Seeking Choice Mean (SD)	-							
Usefulness Mean (SD)	0.40*** (0.28)	-						
Thought Frequency (SD)	0.16*** (0.21)	0.21*** (0.24)	-					
Feelings to Know Mean (SD)	0.26*** (0.27)	0.38*** (0.36)	0.22*** (0.25)	-				
Feelings Never to Know Mean (SD)	-0.22*** (0.24)	-0.30*** (0.33)	-0.19*** (0.26)	-0.25*** (0.36)	-			
Feelings to Know - Feelings Never to Know Mean (SD)	0.29*** (0.26)	0.41*** (0.33)	0.26*** (0.26)	0.84*** (0.13)	-0.71*** (0.23)	-		
Expectations Mean (SD)	0.00 (0.18)	0.02 (0.22)	-0.27*** (0.23)	0.02 (0.22)	0.09*** (0.21)	-0.04 (0.21)	-	
Confidence in estimation Mean (SD)	-0.02 (0.21)	-0.01 (0.25)	0.11**** (0.24)	0.06* (0.26)	0.04 (0.21)	0.02 (0.25)	0.04 (0.28)	-

Information-Seeking Choice Rating: -3 ('*definitely don't want to know*') to +3 ('*definitely want to know*'); **Usefulness:** -3 ('*not useful*') to +3 (*very useful*); **Thought Frequency:** -3 ('*never*') to +3 ('very *often*'); **Feelings to Know:** -3 ('*very bad')* to +3 ('*very good*'); **Feelings Not to Know:** -3 ('*very bad')* to (+3 '*very good*'); **Expectations:** -3 (*'not at all this trait')* to +3 (*'very much this trait';* scores were reversed for negative valanced stimuli); **Confidence in Estimation:** -3 ('*not certain')* to +3 (*'very certain'*).

Supplementary Study (Study 5).

Testing additional ratings. We ran a fifth study (N= 48) to examine if additional ratings explain information-seeking when competing for variance with the ratings described in the main text. Forty-five participants who passed the attention checks and for whom data variability allowed the generation of all beta coefficients were included in the analysis.

The procedure was exactly as in Study 1 except that three additional ratings were included:

(1) **Distinctiveness.** One may hypothesise that individuals would be more interested in receiving information about traits they believe make them unique relative to others. We thus asked subjects to rate 'how much do you differ from others on *this trait*?' on a scale from -3 (Not Different) to +3 (Very Different)
(2) **Sense Making.** One may hypothesise that individuals would be more interested in receiving information that will help them make sense of things that happened in their lives. We thus asked subjects 'if you knew how others rated you on *this trait*, would it help you make sense of things that happened in your life?' on a scale from -3 (Not at All) to +3 (Very Much)
(3) **Recency.** One may hypothesise that individuals would be more interested in receiving information about topics they had contemplated lately. We thus asked subjects 'before today, when was the last time you thought of whether others view you on *this trait*?' on a scale from -3 (Never) to +3 (Last 24-Hours).

We then entered the three ratings above into a linear mixed effect model predicting information-seeking choice along with the three questions from our hypothesised model as well as participants' confidence in their expectations of what the information would reveal. A model with random effects to item and subject did not converge. In line with recommendations (Barr et al., 2013), we reduced the random effect structure until the model was able to converge. This occurred when random effects were assigned to subject and not item. The results revealed that participants were more likely to seek information for topics they thought of often (Cognitive Utility $\beta = 0.102 \pm 0.032$ (SE), $t(65.84) = 3.228$, $p = 0.002$), when they expected to feel better after knowing than not knowing (Hedonic Utility $\beta = 0.116 \pm 0.044$ (SE), $t(31.78) = 2.625$, $p = 0.013$), as well as tended to want information more when they expected information to be useful (Instrumental Utility $\beta = 0.071 \pm 0.039$ (SE), $t(23.41) = 1.810$, $p = 0.08$) and for topics they thought off recently (Recency: $\beta = 0.084 \pm 0.047$ (SE), $t(34.71) = 1.786$, $p = 0.08$). Confidence ($\beta = 0.045 \pm 0.046$ (SE), $t(37.34) = 0.974$, $p = 0.336$), Distinctiveness ($\beta = 0.033 \pm 0.034$ (SE), $t(48.80) = 0.950$, $p = 0.347$) and Sense Making ($\beta = -0.018 \pm 0.045$ (SE), $t(36.69) = -0.389$, $p = 0.7$) were not significant predictors (**Supplementary Figure 2.1**).

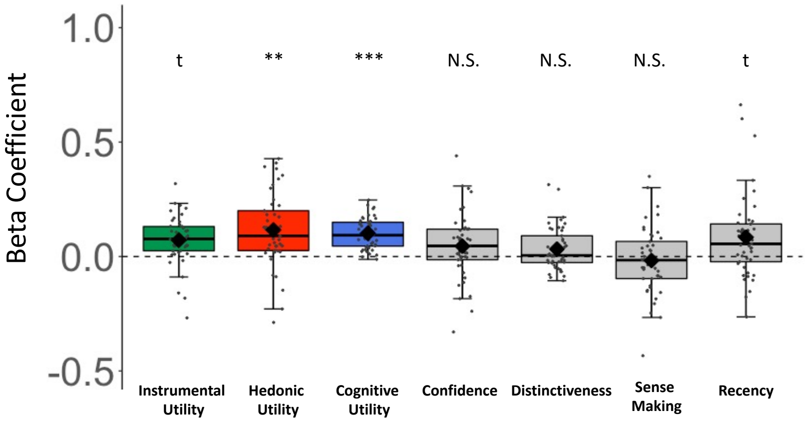

Supplementary Figure 2.1: Information-Seeking Motives. Plotted are beta coefficients from a linear mixed effects model predicting information-seeking (N = 45 subjects), which shows participants want information more when they expect information to make them feel better than ignorance (Hedonic Utility), and for topics they think of often (Cognitive Utility), as well as tend to want information more when its Instrumental Utility is high and for topics they thought of recently (Recency). The horizontal lines indicate median values, boxes indicate 25–75% interquartile range and whiskers indicate 1.5 × interquartile range; individual scores are shown as dots. Distinctiveness (i.e., how much one differs from others on a trait), Sense Making (i.e., whether knowing would help make sense of things that happened in one's life), and Confidence, were not significant predictors of information-seeking. *** = P <0.001.**P <0.01, t = trend, N.S. = not significant (two-sided).

Information-seeking Task

Experiments 1 & 2

(a) Block 1: Measure of information-seeking (40 trials)

Would you like to know how your family/friends rated you on being: **Kind**?

0	0	0	0	0	0
Definitely Don't Want to Know	Don't Want to Know	Somewhat Don't Want to Know	Somewhat Want to Know	Want to Know	Definitely Want to Know

Would you like to know how your family/friends rated you on being: **Mean**?

0	0	0	0	0	0
Definitely Don't Want to Know	Don't Want to Know	Somewhat Don't Want to Know	Somewhat Want to Know	Want to Know	Definitely Want to Know

(b) Block 2: Measure of information-seeking (40 trials)

With regard to being **Kind:**

What rating do you think your family and friends will give you?

0	0	0	0	0	0	0
Not at all this trait (-3)	(-2)	(-1)	(0)	(1)	(2)	Very much this trait (3)

How **CERTAIN** are you about your estimate?

0	0	0	0	0	0	0
Not Certain (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Certain (3)

How **USEFUL** would it be to know how your friends/family have rated you?

0	0	0	0	0	0	0
Not Useful (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Useful (3)

How would you **FEEL** if you got to find out how you were rated?

0	0	0	0	0	0	0
Very Bad (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Good (3)

How would you **FEEL** if you **NEVER** get to find out how you were rated?

0	0	0	0	0	0	0
Very Bad (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Good (3)

How **OFTEN** do you think about **Kindness**?

0	0	0	0	0	0	0
Never (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Often (3)

Experiment 3

(c) Block 1: Measure of information-seeking (40 trials)

Do you want to know what the Gross Domestic Profit is?

0	0	0	0	0	0
Definitely Don't Want to Know	Don't Want to Know	Somewhat Don't Want to Know	Somewhat Want to Know	Want to Know	Definitely Want to Know

Do you want to know what your health expenses are?

0	0	0	0	0	0
Definitely Don't Want to Know	Don't Want to Know	Somewhat Don't Want to Know	Somewhat Want to Know	Want to Know	Definitely Want to Know

(d) Block 2: Measure of information-seeking (40 trials)

For the question about: **what the Gross Domestic Profit is?**

What do you think the answer is?

0	0	0
Low	Average	High

How **CERTAIN** are you in your answer?

0	0	0	0	0	0	0
Not Certain (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Certain (3)

How **USEFUL** would it be to know the answer to this question?

0	0	0	0	0	0	0
Not Useful (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Useful (3)

How would you **FEEL** if you got to find out the answer?

0	0	0	0	0	0	0
Very Bad (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Good (3)

How would you **FEEL** if you **NEVER** get to find out the answer?

0	0	0	0	0	0	0
Very Bad (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Good (3)

How **OFTEN** do you think about **Gross Domestic Profit**?

0	0	0	0	0	0	0
Never (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Often (3)

Experiment 4

(e) Block 1: Measure of information-seeking (40 trials)

Would you like to know if you have a gene that increases your likelihood of: **Alzheimer's disease?**?

0	0	0	0	0	0
Definitely Don't Want to Know	Don't Want to Know	Somewhat Don't Want to Know	Somewhat Want to Know	Want to Know	Definitely Want to Know

Would you like to know if you have a gene that increases your likelihood of: **Good Memory?**

0	0	0	0	0	0
Definitely Don't Want to Know	Don't Want to Know	Somewhat Don't Want to Know	Somewhat Want to Know	Want to Know	Definitely Want to Know

(f) Block 2: Measure of information-seeking (40 trials)

With regard to the gene that increases the likelihood for: **Alzheimer's disease**

How **LIKELY** is it that you carry this gene?

0	0	0	0	0	0	0
Not Likely (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Likely (3)

How **CERTAIN** are you about your estimate?

0	0	0	0	0	0	0
Not Certain (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Certain (3)

How **USEFUL** would it be to know whether or not you carry this gene?

0	0	0	0	0	0	0
Not Useful (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Useful (3)

How would you **FEEL** if you got to find out whether or not you carry this gene?

0	0	0	0	0	0	0
Very Bad (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Good (3)

How would you **FEEL** if you **NEVER** get to find out whether or not you carry this gene?

0	0	0	0	0	0	0
Very Bad (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Good (3)

How **OFTEN** do you think about: **Alzheimer's disease?**

0	0	0	0	0	0	0
Never (-3)	(-2)	(-1)	(0)	(1)	(2)	Very Often (3)

Supplementary Figure 2.2: Information-Seeking Task. (a) For Studies 1 & 2, participants were asked to imagine that their family/friends have rated them on 40 different attributes. They then indicated whether they would like to know how they have been rated (e.g., on being kind) from '*definitely don't want to know'* to '*definitely want to know'.* **(c)** For Study 3, participants were asked to indicate whether they wanted to know different 40 pieces of information related to finance (e.g., what the Gross Domestic Profit is) from '*definitely don't want to know'* to '*definitely want to know'.* **(e)** For Study 4, participants were asked to imagine that we had information about their genetic makeup and asked whether they wanted to know 40 pieces of information related to Health (*"Would you like to know if you have a gene that increases your likelihood of Alzheimer's disease?"*) from '*definitely don't want to know'* to '*definitely want to know'.* **(b,d)** Next, participants provided the following ratings for each stimulus (self-paced): (i) Their expectations regarding how useful each piece of information would be from '*not useful* 'to '*very useful'*; (ii) How they expect to feel if the rating was revealed to them from '*very bad'* to '*very good'*, and how they expect to feel if the rating was never revealed to them from '*very bad'* to '*very good'*.; (iii) How often they think about the topic in question from '*never* 'to 'very *often'*. Each question is displayed separately for each stimulus. Participants were also asked in Studies 1 & 2 to indicate for each attribute (i) what rating they think their family and friends will give them and (ii) how certain they are in that. In Study 3 they were asked to indicate (i) what they thought the answer was (for each stimulus the scale was different for this question and for some stimuli this was an open-ended question) and (ii) how certain they are in their answer. In Study 4, participants indicated their expectations of how likely it is that they carry the gene (from *-3 'not likely'* to *+3 'very likely', e.g.,* how likely is it that you carry this gene?)**.** Finally, participants indicated their confidence in what they expected the information would reveal (from -3 '*not certain'* to +3 *'very certain'*). In Studies 1 & 2, participants filled in questionnaires assessing mental health following the task.

Supplementary Table 2.7. Studies 1 & 2: Correlations (partial R coefficients) between the weight subjects assign to motives of information-seeking and scores on specific psychopathology questionnaires.

Study 1 Psychopathology Questionnaires	**Instrumental Utility (β_1)**	**Hedonic Utility (β_2)**	**Cognitive Utility (β_3)**
Depression	-0.111	0.083	-0.011
Anxiety	-0.093	0.137	-0.197
Apathy	-0.089	-0.037	-0.146
OCD	-0.057	0.068	-0.115
Social Anxiety	0.086	0.116	-0.276*
Alcohol Use Disorder	0.024	-0.082	-0.113
Impulsivity	-0.040	-0.047	0.010
Schizotypy	-0.229*	0.080	-0.108
Eating Disorder	-0.274*	0.087	-0.146
Study 2[1] Psychopathology Questionnaires	**Instrumental Utility (β_1)**	**Hedonic Utility (β_2)**	**Cognitive Utility (β_3)**
Depression	0.088	0.249***	-0.138
Anxiety	0.086	0.209*	-0.132
Apathy	0.061	0.236**	-0.269***
OCD	0.078	-0.012	-0.177*
Social Anxiety	0.088	0.190*	-0.204
Alcohol Use Disorder	0.169	-0.083	-0.065
Impulsivity	0.058	-0.061	-0.355***
Schizotypy	0.106	0.197*	-0.263***
Eating Disorder	0.050	-0.059	-0.177

*Bonferroni corrected: ***p< 0.006, No correction: **p<0.01, *p<0.05* (two-sided).
Displayed are the partial R coefficients controlling for gender and age.
The following questionnaires were used to assess psychopathology: **Depression** = Self-Rating Depression Scale; **Anxiety** = State-Trait Anxiety Inventory; **Apathy** = Apathy Evaluation Scale; **OCD** = Obsessive-Compulsive Inventory – Revised; **Social Anxiety** = Liebowitz Social Anxiety Scale; **Alcohol Use Disorder** = Alcohol Use Disorder Identification Test; **Impulsivity** = Barratt Impulsivity Scale; **Schizotypy** = Short Scales for Measuring Schizotypy; **Eating Disorder** = EAT-26.
[1]Data displayed for Study 2 are psychopathology scores correlated with the mean beta coefficient for each motive over Time 1 and Time 2.

Supplementary Table 2.8. *Studies 1 - 4: AIC scores for all models.*

Model (AIC Score)	Study 1	Study 2, Time 1	Study 2, Time 2	Study 3, Time 1	Study 3, Time 2	Study 4
Instrumental + Hedonic + Cognitive	**11066.38**	**24863.27**	**17242.26**	**17595.42**	**11625.84**	**15456.62**
Instrumental + Hedonic + Cognitive + Confidence	11070.74	24869.11	17247.84	17608.72	11639.83	15463.2
Hedonic + Cognitive + Confidence	11107.96	24907.92	17260.81	17874.96	11922.91	15842.82
Hedonic + Cognitive	11109.39	24905.13	17257.45	17885.18	11911.51	15834.96
Instrumental + Cognitive + Confidence	11122.51	25020.58	17342.45	17663.99	11663.03	15568.16
Instrumental +Hedonic + Confidence	11186.01	24914.37	17303.71	17707.74	11756.56	15523.26
Instrumental + Cognitive	11301.75	25067.15	17405.32	17776	11801.1	15654.25
Instrumental + Hedonic	11183.4	24915.38	17301.11	17750.06	11750.34	15517.1
Cognitive + Confidence	11200.15	25148.86	17410.02	18090.35	12062.3	16135.68
Cognitive	11215.41	25157.28	17404.61	18078.92	12051.43	16134.31
Hedonic + Confidence	11247.08	24977.47	17329.6	18107.07	12179.93	15966.8
Hedonic	11251.14	24982.45	17329.49	18111.19	12185.06	15957.91
Instrumental + Confidence	11301.75	25067.15	17405.32	17776	11801.1	15654.25
Instrumental	11312.66	25082.81	17406.1	17768.03	11794.75	15640.29

Supplementary Table 2.9. *Studies 1 – 4: R^2 for hypothesised model.*

Variance Test	Study 1	Study 2, Time 1	Study 2, Time 2	Study 3, Time 1	Study 3, Time 2	Study 4
Conditional R^2	0.568	0. 429	0. 345	0. 422	0. 515	0. 523
Marginal R^2	0. 035	0. 021	0. 035	0. 134	0. 171	0. 089
Participants' Mean R^2	0.234	0.174	0.175	0.294	0.351	0.312

Conditional and Marginal R^2 are calculated for mixed models (Nakagawa & Schielzeth, 2013). Conditional R^2 reflects variability explained by the full model (random and fixed effects). Marginal R^2 reflects the variability explained only by fixed effects, not considering random effects. This is less informative in our case, because of the large individual differences in the weight subjects assign to the different motives, which are captured by random slopes. Participants' mean R^2 represents the average R^2 calculated for each subjects' linear model separately.

Study 1 & 2, Stimuli

All stimuli for Studies 1, 2, and 5 (traits) were adapted from Allport and Odbert's (1936) trait-word list. Negative stimuli for Study 4 (health conditions) were adapted from a WHO report indicating common causes of death. Positive stimuli for Study 4 (health conditions) and all stimuli for Study 3 (finance questions) were developed by the authors.

Study 3, Time 1 Stimuli

The following sentences followed the question *"Do you want to know…"*

what your health expenses are?
whether your credit card information has been compromised?
what the Gross Domestic Profit is?
in which percentile income bracket you fall into in your country?
how your salary compares to others doing a similar job?
how much more/less individuals of your gender make in your job relative to the opposite gender?
how much more/less individuals of your race make in your job relative to a different race?
what the gender ratio is for individuals with your job title?
how the Dollar will compare to the Euro at the end of 2020?
how the Dollar will compare to the Yen at the end of 2020?
how much your property is worth?
how much you will pay in travel expenses this year?
how much your phone bill will be this year?
what the unemployment rate is in your country?
what the unemployment rate is in Australia?
what you had spent on dining this year?
the value of gold?
what the financial impact of Brexit is on the global economy?
what your bank balance will be on December 31st, 2020?
what your yearly income will be 5 years from now?
how much each year you will receive from your pension when you retire?
how the stock market will be performing 1 year from now?
whether you will change professions in the future?
what age you will retire?
where you will live 5 years from now?
what your credit score will be in 5 years?
how much your grandparents made when they were younger?
how much your next vacation will cost?
which phone carrier provides the best deal?
whether a family member will get promoted?
whether a family member needs a loan?
how much in total you will spend on utilities in 2021?
how much you will be required to pay in federal and state taxes next year?

what the price of oil and gas will be in 5 years?
the earnings of your favourite celebrity for 2020?
the average value of homes in your neighbourhood?
if you can make more/less income on different platforms from the one you are on now?
what the top financial investments advisors are recommending?
the value of the Dow Jones?
*what you had spent on clothing this year***?**

Study 3, Time 2 Stimuli

The following sentences followed the question *"Do you want to know…":*

what Warren Buffet recommends investing in?
the earnings of Donald Trump for 2020?
the value of the NASDAQ?
how much Apple is worth?
how much your hairdresser makes?
how much your water and gas bill will be this year?
what the unemployment rate is in Europe?
what the unemployment rate is in Asia?
which bank gets the best user ratings?
the average cost of a 2-carat diamond?
what the financial impact of Covid-19 is on the global economy?
which airline provides the best deal?
the average value of homes in your city?
if you can make more/less in a different job right now?
how your salary compares to others doing a similar job in a different country?
how much more/less individuals of your race make on average compared to those of different race?
what the gender ratio is for individuals in the bottom 1% of earners?
what the gender ratio is for individuals in the top 1% of earners?
how the US Dollar will compare to the British Pound at the end of 2020?
how the US Dollar will compare to the Canadian Dollar at the end of 2020?
how much your property will be worth in 5 years?
the mortgage rates right now?
how much your assets will be worth in 5 years?
what your yearly income will be next year?
when the next income relief stimulus package will be delivered?
what the inflation rate will be 1 year from now?
whether you will move States for work in the future?
how much you will receive upon retirement?
how large your house will be 5 years from now?
if you will have debt in 5 years?
what your daily expenses are?
how much your next vehicle will cost?
whether a close friend will get promoted?

whether a close friend needs a loan?
how much in total you will spend on commuting in 2021?
whether your bank details have been compromised?
what the national debt is?
how much you will pay for Social Security?
what the price of silver will be in 5 years?
in which percentile income bracket you will fall in 5 years?

Study 4, Stimuli

The following sentences followed the question *"Would you want to know if you have a gene that increases your likelihood of…":*

Alzheimer's Disease?
A Youthful Appearance?
Dementia?
Diabetes?
Good Concentration?
Arthritis?
Stroke?
Infertility?
Clear Skin?
Lactose Intolerance?
Strong Immune System?
Liver Disease?
Obesity?
OCD?
Parkinson's Disease?
Long Life Expectancy?
Prostate/Breast Cancer?
Schizophrenia?
Good Hand and Eye Coordination?
Skin Cancer?
Serious Covid-19 Symptoms?
Healthy Cholesterol Level?
Healthy Sleep Cycles?
Brain Tumor?
Fresh Breath?
Good Memory?
Depression?
High Fertility?
High Lung Capacity?
Good Vision?
Heart Disease?
High Intelligence/IQ?
Fast Metabolism?
Leukemia?
Strong Joints?

Being Athletic?
Lung Cancer?
High Tolerance to Stress?
Strong Bones?
Sexual Dysfunction?

Information Avoidance (Study 1-4)

Participants indicated they would rather avoid knowledge (that is selected -3, -2, or -1 on the information-seeking question) on 37.83% of the trials in Study 1, 23.4% in Study 2 Time 1, 30.2% in Study 2 Time 2, 23.9% in Study 3 Time 1, 28.8% in Study 3 Time 2, and 23.5% in Study 4.

A paired-samples t-test was conducted to compare the mean ratings on the different scales when individual indicated they preferred knowledge (+3, +2, +1) compared to when they indicated they preferred ignorance (-3, -2, -1). The results are presented in **Supplementary Table 2.10** below.

Supplementary Table 2.10. Paired-samples t-test comparing the mean ratings on the different scales between trials on which participants selected knowledge and ones in which they selected to avoid knowledge.

Study 1	**Mean Difference Knowledge Trials minus Avoidance Trials (SD)**	**t**	**df**	**p**
Usefulness	0.57 (1.04)	4.450	65	0.0001
Thought Frequency	0.60 (0.97)	5.005	65	0.0001
Feelings to Know	0.70 (0.95)	5.919	65	0.0001
Feelings Never to Know	-0.35 (0.74)	-3.869	65	0.0001
Expected rating of others	0.91 (1.26)	5.862	65	0.0001
Confidence in estimation	0.08 (0.88)	0.778	65	0.439
Study 2, Time 1				

	Mean Difference Knowledge Trials minus Avoidance Trials (SD)	t	df	p
Usefulness	0.59 (1.02)	-6.814	137	0.0001
Thought Frequency	0.37 (1.00)	-4.335	137	0.0001
Feelings to Know	0.63 (0.78)	-9.472	137	0.0001
Feelings Never to Know	-0.29 (0.62)	5.583	137	0.0001
Expected rating of others	0.75 (1.04)	-8.502	137	0.0001
Confidence in estimation	0.19 (0.70)	-3.198	137	0.002

Study 2, Time 2	Mean Difference Knowledge Trials minus Avoidance Trials (SD)	t	df	p
Usefulness	0.56 (0.76)	7.796	110	0.0001
Thought Frequency	0.46 (0.80)	6.015	110	0.0001
Feelings to Know	0.56 (0.75)	7.871	110	0.0001
Feelings Never to	-0.26 (0.56)	-4.938	110	0.0001
Expected rating of others	0.96 (1.12)	9.105	110	0.0001
Confidence in estimation	0.06 (0.65)	0.953	110	0.343

Study 3, Time 1	Mean Difference Knowledge Trials minus	t	df	p

	Avoidance Trials (SD)			
Usefulness	1.60 (1.33)	12.470	107	0.0001
Thought Frequency	0.75 (0.82)	9.465	106	0.0001
Feelings to Know	1.13 (1.13)	10.368	107	0.0001
Feelings Never to Know	-0.54 (0.59)	-9.606	106	0.0001
Confidence in estimation	0.46 (0.91)	5.264	107	0.0001

Study 3, Time 2	**Mean Difference Knowledge Trials minus Avoidance Trials (SD)**	**t**	**df**	**p**
Usefulness	1.60 (1.29)	10.758	74	0.0001
Thought Frequency	1.39 (1.17)	10.298	74	0.0001
Feelings to Know	0.80 (0.70	9.921	74	0.0001
Feelings Never to Know	-0.65 (0.77)	-7.253	74	0.0001
Confidence in estimation	0.37 (0.83)	3.884	74	0.0001

Study 4	Mean Difference Knowledge Trials minus Avoidance Trials (SD)	t	df	p
Usefulness	1.59 (1.34)	11.219	89	0.0001
Thought Frequency	0.73 (1.06)	6.520	89	0.0001
Feelings to Know	0.70 (0.88)	7.604	89	0.0001
Feelings Never to Know	-0.48 (0.77)	-5.954	89	0.0001
Expectations	-0.22 (0.95)	-2.192	89	0.031
Confidence in estimation	0.07 (0.90)	0.750	89	0.455

Information-seeking Choice and Expectation

Below we plot the distribution of participants' ratings on the question of whether they wanted information in Studies 1, 2, 3 ,4 and 5. We also plot the distribution of participants' raw ratings on what they thought the information will reveal in Studies 1, 2, 4 and 5. In Study 3 expectations could not be quantified and high/low numbers do not indicate positive/negative expectations. In Study 1 and 2, participants' who had more negative expectations of what information was to reveal scored higher on the psychopathology factors (Study 1: Anxious-Depression: $r(71) = -0.672$, $p = 0.0001$, Compulsive Behaviour and Intrusive Thought: $r(71) = -0.333$ $p = 0.004$, Social-Withdrawal: $r(124) = -0.571$, $p = 0.0001$; Study 2: Anxious-Depression: $r(124) = -0.519$, $p = 0.0001$, Compulsive Behaviour and Intrusive Thought: $r(124) = -0.360$, $p = 0.0001$, Social-Withdrawal: $r(124) = -0.447$, $p = 0.0001$). Information-seeking choice was not significantly correlated with any of the three psychopathology factors in Study 1 (Anxious-Depression: $r(80) = -0.094$, $p = 0.434$; Compulsive Behaviour and Intrusive Thought: $r(71) = 0.184$, $p = 0.124$; Social-Withdrawal: $r(71) = -0.051$, $p = 0.672$). In Study 2, people who were less likely to want information scored higher on Social-Withdrawal ($r(124) = -0.201$, $p = 0.025$), but not with either Anxious-Depression ($r(124) = -0.007$, $p = 0.939$) nor Compulsive Behaviour and Intrusive Thought ($r(124) = -0.053$, $p = 0.562$).

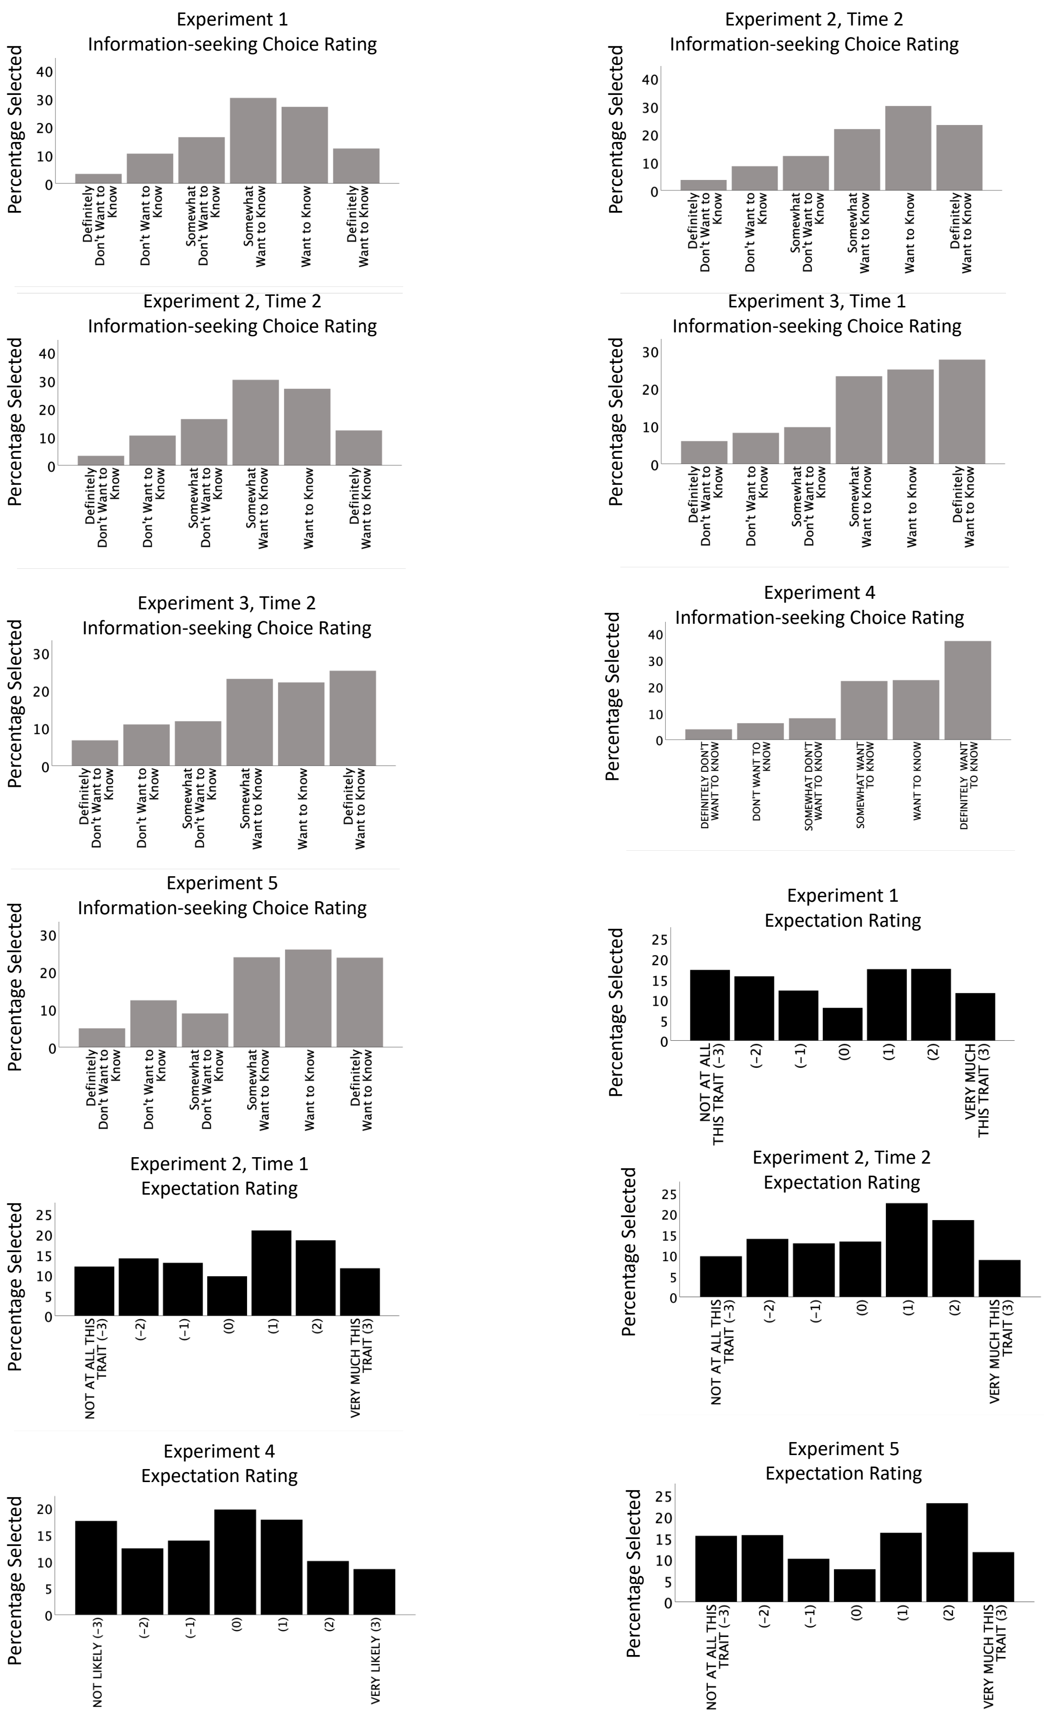

Supplementary Figure 2.3. Plotted are the distributions of participants' ratings on the question of whether they wanted information (grey) in Studies 1, 2, 3, 4, and 5 and participants' ratings on what they thought the information will reveal (black) in Studies 1, 2, 4 and 5.

Supplementary Table 2.11. Studies 1, 2 & 4: Substituting Hedonic Utility for the mean of Hedonic Utility and Rating of Expected Information in the theorised model (provided as response to a reviewer's request).

	B (SE)	df	t-value	p (two-sided)
Study 1				
Instrumental Utility	0.106 (0.03)	55.35	3.800	0.0004
Mean of Hedonic Utility & Expectations	0.175 (0.03)	70.99	6.086	0.0001
Cognitive Utility	0.095 (0.03)	91.79	3.011	0.0003
Study 2, Time 1				
Instrumental Utility	0.076 (0.02)	160.29	4.360	0.0001
Mean of Hedonic Utility & Expectations	0.180 (0.02)	165.02	8.112	0.0001
Cognitive Utility	0.057 (0.02)	171.52	3.729	0.0003
Study 2, Time 2				
Instrumental Utility	0.084 (0.02)	79.44	4.311	0.0001
Mean of Hedonic Utility & Expectations	0.216 (0.03)	108.95	8.318	0.0001
Cognitive Utility	0.092 (0.02)	127.05	4.963	0.0001
Study 4				
Instrumental Utility	0.248 (0.03)	123.27	9.346	0.0001
Mean of Hedonic Utility & Expectations	0.080 (0.02)	121.52	3.435	0.0008
Cognitive Utility	0.115 (0.02)	135.16	7.639	0.0001

Note. Raw scores for Expectations were reverse scored for negative stimuli. Raw scores were then centred within participants and entered into a linear mixed effect model with a random intercept and slope for subject and random intercept for item.

Supplementary Table 2.12. Studies 1 & 2: Correlation between transdiagnostic factors and information-seeking motives.

	Instrumental Utility	Hedonic Utility	Cognitive Utility
Study 1			
Anxious Depression	-0.146	0.063	-0.132
Compulsive Behaviour & Intrusive Thought	-0.193	0.077	-0.181
Social Withdrawal	0.017	0.075	-0.274*
Study 2			
Anxious Depression	0.072	0.243**	-0.197*
Compulsive Behaviour & Intrusive Thought	0.116	-0.046	-0.214*
Social Withdrawal	0.076	0.226*	-0.153
Note. **Controlling for age and gender.** ***p<0.01, *p<0.05*** **(two-sided).**			

Chapter 3

Replication of results using the Hu and Liu Valence lexicon (2004).

To validate the findings in the main article, we implemented the exact same approach described within but this time with a different valence lexicon (Hu & Liu, 2004), which categorises 2006 words as positive and 4783 as negative. Other than that, the score for positive words and negative words were calculated exactly as described in the main text. Note, the Hu and Liu lexicon does not categorise specific emotions, so analysis was only conducted using its Negative and Positive scores.

Valence of webpages provides a marker of mental health. We again conducted two separate mixed ANOVA's for the Positive and Negative score of webpages visited, but using the Hu and Liu Valence Lexicon (described above). In the first mixed ANOVA, psychopathology scores ('Anxious-Depression', 'Social-Withdrawal', 'Compulsive-Behaviour and Intrusive Thought') was indicated as a within-subjects factor and the Negative Valence score of the webpages that participants browsed (Z-scored) was input as a within-subject modulating factor. Participants' age and gender were also indicated between-subject modulating covariates (both Z-scored). We observed a significant main effect of the Negative Score of webpages that participants browsed on psychopathology scores (**Study 1:** $F(1,284) = 4.083$, $p = 0.044$, partial eta square = 0.014; **Study 2:** $F(1,442) = 6.462$, $p = 0.011$, partial eta square = 0.014). The second mixed ANOVA was identical to the first except that the Positive Valence score was input as a within-subject modulating factor instead of the Negative score. We did not observe a significant main effect of the Positive score of webpages that participants browsed on psychopathology scores (**Study 1:** $F(1,284) = 0.155$, $p = 0.695$ partial eta square = 0.001; **Study 2:** $F(1,442) = 4.168$, $p = 0.042$, partial eta square = 0.009). Together, these results indicate that Negative valence of webpages that people expose themselves to online is indicative of their mental health.

Mood is bi-directionally related to browsing negatively valenced webpages. We again tested whether participants pre-browsing mood was related to the valence of information they browsed but using the Hu and Liu Valence Lexicon (described above). To do this, in Study 1 we ran two separate mixed effect models each including participants pre-browsing mood ratings as fixed and random effects along with age and gender as fixed effect predicting the Negative score and Positive score of webpages visited separately. In Study 2, as we only have 1 observation per participant for each variable of interest, (compared to 5 in Study 1), we ran two simple linear regressions predicting the Negative score and Positive score, separately, from pre-browsing mood ratings, controlling for age and gender. Participants pre-browsing mood was associated with the Negative score of webpages visited (**Study 1:** $\beta = -0.071 \pm 0.041$ (SE), $t(364.49) = -1.720$, trend $p = 0.086$; **Study 2:** $\beta = -0.001 \pm 0.000$ (SE), $t(399) = -2.293$, $p = 0.022$). With regard to participants pre-browsing mood predicting the Positive score of webpages visited, we once again did not observe a significant effect (**Study 1:** $\beta = -0.029 \pm 0.041$ (SE), $t(94.11) = -0.723$, $p = 0.471$; **Study 2:** $\beta = 0.000 \pm 0.001$ (SE), $t(399) = 0.861$, $p = 0.390$).

Next, we tested whether the valence of the webpages that participants browsed had an impact on their mood directly after browsing the internet. To test this, in Study 1, we once again ran two mixed effect models, each predicting post browsing mood ratings from either the Hu and Liu (2004) Negative or Positive score of webpages visited (input as a fixed and random effect). Both models included participants pre-browsing mood ratings as a fixed and random effect, along with age and gender as a fixed effect. In Study 2, we ran two simple linear regressions, both predicting participants post browsing mood ratings from either the Negative or Positive score of webpages visited. Both models included participants pre-browsing mood ratings, age and gender as control variables. We observed that Negative score of webpages visited was related to participants post browsing mood ratings controlling for pre-browsing mood, age and gender (**Study 1:** $\beta = -0.035 \pm 0.018$ (SE), $t(715.58) = -1.977$, $p = 0.048$; **Study 2:** $\beta = -1.413 \pm 0.630$ (SE), $t(399) = -2.244$, $p = 0.025$). In particular, participants expressed worse mood post browsing when they browsed more negatively valenced webpages. We also observed a significant effect of the Positive score of webpages visited on participants post browsing mood ratings using the Hu and Liu Valence Lexicon (2004), which we did not see when using the NRC lexicon (Mohammad, 2018): (**Study 1:** $\beta = 0.036 \pm 0.018$ (SE), $t(604.29) = 1.990$, $p = 0.047$; **Study 2:** $\beta = 1.670 \pm 0.626$ (SE), $t(399) = 2.669$, $p = 0.008$).

Supplementary Table 3.1. *Model results for Age and Gender (i.e., controls).*

Full Model	Page ß	Gender ß
Study 1: (Mean Psychopathology ~ Negative score + Age + Gender)	-0.179*** (i.e., younger participants report more psychopathology symptoms).	0.132** (i.e., females report more psychopathology symptoms).
Study 2: (Mean Psychopathology ~ Negative score + Age + Gender)	-0.135*** (i.e., younger participants report more psychopathology symptoms).	0.122*** (i.e., females report more psychopathology symptoms).
Study 1: (Negative score ~ Pre-Mood + Age + Gender)	0.008* (i.e., older participants browse more negative webpages).	-0.141(not significant).
Study 2: (Negative score ~ Pre-Mood + Age + Gender)	-0.071 (not significant).	-0.177*** (i.e., males browse more negative webpages).
Study 1: (Post Mood ~ Negative score + Pre-Mood + Age + Gender)	0.000 (not significant).	-0.080** (i.e., males report better mood after browsing the web).
Study 2: (Post Mood ~ Negative score + Pre-Mood + Age + Gender)	-0.029 (not significant).	-0.070* (i.e., males report better mood after browsing the web).
*** = P <0.001, ** = P <0.01 * = P <0.05		

Fear sentiment of webpages browsed associated with mental health. To test whether the specific emotions of webpages browsed were associated with mental health we quantified the percentage of *Anger, Fear, Anticipation, Trust, Surprise, Sadness, Joy* and *Disgust* associated words (as defined by the NRC Emotion Lexicon; Mohammad and Turney, 2013) out of all words on each webpage participants browsed. For each day separately, we then calculated the average emotion score of the webpages visited by each participant and then averaged these scores across the five days. We then input the eight Emotion scores, along with age and gender into a stepwise regression predicting the mean of the three psychopathology factors (i.e., '*Anxious-Depression*', '*Social-Withdrawal*' and '*Compulsive-Behaviour and Intrusive Thought*'). A stepwise regression was ideal in this case as the variance inflation factor (VIF) of some predictor variables was high (e.g., greater than 3). The winning model included the Fear score of webpages (**Study 1:** $\beta = 0.105 \pm 0.042$ (SE), $t(288) = 2.500$, $p = 0.013$, $r = 0.146$; **Study 2:** $\beta = 0.087 \pm 0.034$ (SE), $t(446) = 2.569$, $p = 0.011$, $r = 0.121$) as well as age and gender, suggesting that those with poorer mental health browse more fear related webpages. The Intraclass Correlation Coefficient (ICC) of the Fear scores across the 5 days revealed statistically significant moderate stability (ICC = 0.505, $p < 0.001$), indicating that the tendency to consume text high in fear words is likely due both to trait and state.

Supplementary References

Allport, G. W., & Odbert, H. S. (1936). Trait-names: A psycho-lexical study. Psychological monographs, 47(1), i.

Hu, M., & Liu, B. (2004). Mining and summarizing customer reviews. In Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining (pp. 168-177).

Mohammad, S. M., & Turney, P. D. (2013). Crowdsourcing a word–emotion association lexicon. Computational intelligence, 29(3), 436-465.

Nakagawa, S., & Schielzeth, H. (2013). A general and simple method for obtaining R2 from generalized linear mixed-effects models. Methods in ecology and evolution, 4(2), 133-142.