

Is There a Reverse Privacy Paradox? An Exploratory Analysis of Gaps Between Privacy Perspectives and Privacy-Seeking Behaviors

Jessica Colnago
Google
Sunnyvale, California, USA
jcolnago@google.com

Lorrie Faith Cranor
Carnegie Mellon University
Pittsburgh, Pennsylvania
lorrie@cmu.edu

Alessandro Acquisti
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
acquisti@andrew.cmu.edu

ABSTRACT

Privacy scholars have long studied, and argued about, a so-called privacy paradox—the alleged gap between individuals’ claims of caring about privacy and their actual behaviors. This manuscript explores whether a different type of mismatch occurs in an online sample of US participants: a mismatch between participants’ dismissive perspectives on privacy and their privacy-protective behaviors. In a series of online studies with Prolific US participants we tackle two research questions: is there evidence of mismatches between (dismissive) privacy perspectives, and (protective) privacy behaviors? If so, what can explain those mismatches? In a Behavior Elicitation study, we collect a corpus of privacy-regulating and privacy-protective behaviors. Next, in Study 1, we find evidence that engagement in a broad array of privacy behaviors is, in fact, very common in our sample. We also find that mismatches between dismissive privacy perspectives and protective behaviors emerge in a large proportion of participants. Finally, in Study 2, we uncover several common but distinct reasons for those mismatches, including construing seemingly protective behaviors as motivated by reasons other than privacy, and nuanced stances on when to express privacy concern. Collectively, the results indicate that individuals who are seemingly dismissive of privacy concerns engage in behaviors that can be construed as privacy-seeking. The findings highlight the nuances of individual privacy decision-making and suggest that public policy related to privacy should account for the evidence for widespread privacy-seeking behaviors.

KEYWORDS

Privacy, Altman, Human Behavior, Privacy Paradox

1 INTRODUCTION

In 2009, following a speech by Supreme Court Justice Antonin Scalia that seemingly belittled online privacy concerns, Professor Joel Reidenberg had students engage in an unusual project. He asked them to create a dossier about Justice Scalia entirely based on what could be found on the Internet. The resulting document included “the justice’s home address and home phone number, his wife’s personal e-mail address and the TV shows and food he prefers.” Justice Scalia was not amused. The exercise, he said in a

statement, was “an example of perfectly legal, abominably poor judgment” [14].

Is it common to say that societal privacy is not that important, only to feel violated when *ours* is intruded upon? Are reactions such as Scalia’s unique or commonplace? And what about behaviors such as those of Facebook’s CEO Mark Zuckerberg—who, after publicly wondering whether privacy was a social norm of the past [30], sought seclusion for his Hawaii property by building a wall around it [27]?

For many years scholars have studied and argued around a so-called privacy paradox—an alleged gap, or mismatch, between individuals’ claims of caring about privacy and their actual behaviors. In this manuscript we explore whether a different type of mismatch occurs. We define this mismatch as a gap between an individual’s claimed privacy perspectives or mental states regarding privacy (such as stated attitudes, preferences, desires, or concerns), interpreted by the observer as dismissive (that is, lacking of interest, care, or concern about privacy), and that individual’s observed or reported behavior, interpreted as privacy seeking, regulating, or protecting.

Our interest in this matter originates from three factors. The first consists of anecdotal observations the authors made of patterns of apparent mismatches emerging in everyday scenarios—such as the above-mentioned story about Justice Scalia and Professor Reidenberg. The second factor is theory- and literature-driven, grounded in Irwin Altman’s research on privacy [10]. Rather than conceiving of privacy merely as a process of seclusion and hiding (or, even more narrowly, as a process of information protection), Altman theorized privacy as a process of boundary regulation, a dynamic and dialectic opening and closing of the self to others. Altman famously observed that privacy behaviors in everyday life are ubiquitous and almost instinctual [10]—for instance, when we lower our voice during the part of a conversation that we do not want third parties to hear. But if so, when individuals claim their life to be an “open book,” and to have “nothing to hide,” do they literally mean that they hold nothing private, and never engage in any of the behaviors Altman considered as boundary-regulating? Or, instead, do they engage in some of those behaviors from time to time but do not always construe them as *privacy* behaviors? A third factor emerged from our own research. While conducting interviews to understand privacy perspectives of a personalized privacy assistant [16], we found participants who, in the early portion of the interview, would claim not to be concerned about privacy but then, in its later portion, would delve into minute details of the complex digital privacy-protecting strategies they regularly engaged in online.

This work is licensed under the Creative Commons Attribution 4.0 International License. To view a copy of this license visit <https://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.



Proceedings on Privacy Enhancing Technologies 2023(1), 455–476
© 2023 Copyright held by the owner/author(s).
<https://doi.org/10.56553/popets-2023-0027>

These factors motivate our investigation and the research questions we tackle in this manuscript: is there general evidence of gaps, or mismatches, between (dismissive) privacy claims and (protective) privacy behaviors (RQ1)? If so, what are the possible explanations for that gap (RQ2)?

If those mismatches do occur, the above discussion already highlighted an array of potential explanations for them. Perhaps, since different individuals construe privacy differently, privacy-embracing and privacy-dismissive claims depend less on whether we actually care for and engage in an Altmanian process of boundary regulation, and more on whether we embrace or reject *other people's* notions of, and value assigned to, privacy protection. After all, privacy boundaries—and even the way we construe them—vary from culture to culture [5] and from person to person [4]. Perhaps, some individuals claim not to be bothered about privacy precisely because they are already taking measures to address their concerns. And perhaps some individuals feel they have to adapt to the modern circumstances of digital publicness, while trying to carve out some private spaces for themselves within those circumstances.

We tackle our research questions in a series of online studies conducted with US participants on the Prolific platform. In a behavior Elicitation study (N=60), we conduct a between-subjects survey of Prolific US participants to generate a list of privacy-protective behaviors participants report to engage in. In Study 1 we tackle RQ1: we investigate possible gaps between self-reported privacy perspectives and actual behaviors using a within-subjects quantitative survey. We do so by asking a different sample of Prolific US participants about 10 behaviors drawn at random from the list of privacy behaviors generated from the Behavior Elicitation Study (Part A; N=255). Three days later, we ask them to respond to 25 general statements about privacy attitudes, concerns, and preferences, followed by statements related to the behaviors they reported engaging in (Part B; N=73). In Study 2 (N=49), we tackle RQ2. We invite participants who, in Study 1, had expressed negative perspectives on privacy (generally or in specific situations) to participate in a new study. Study 2 is a qualitative survey in which we ask participants questions about the mismatches we observed between their stated privacy perspectives and behaviors. The questions aim at understanding the possible reasons behind the detection of a gap between perspectives and behaviors.

The results of Study 1 suggest that engagement in a broad array of privacy behaviors is very common: the vast majority of participants reported engaging in a *majority* of the behaviors that we collected in the Behavior Elicitation Study. They also provide evidence of mismatches (dismissive privacy perspectives, but evidence of privacy behaviors) across a majority of our sample. The results of Study 2 present evidence for multiple, not mutually exclusive explanations for those mismatches, including construing seemingly protective behaviors as motivated by reasons other than privacy, and nuanced stances on when to express privacy concern.

Linking our research to the somewhat contentious literature on the so-called privacy paradox [6, 29, 36] helps us illustrate our contribution and its lessons for public policy. Some of the confusion surrounding the privacy paradox is due to the contrasting implications scholars and observers have taken from it (see Section 2.3). One interpretation of the original privacy paradox, grounded on so-called revealed-preferences economic arguments, is that behaviors

in the market reflect true preferences; hence, if observed behaviors deviate from claimed privacy desires, it must mean that people do not care that much for privacy and policy interventions are therefore unwarranted. A different interpretation (which we subscribe to) is that even privacy-conscious individuals face systemic behavioral and economic hurdles that can make their desired degrees of privacy hard to attain [6]. Those hurdles contribute to explaining the attitudes/behavior gap and boost the importance of policy interventions in the privacy domain. We can learn complementary lessons from the evidence we present here of a reverse gap between dismissive privacy perspectives and actually protective behaviors. First, evidence of a reverse gap expands our understanding of the extent to which even individuals who claim to not be particularly concerned about their privacy end up engaging in behaviors that appear privacy-protective. Second, it provides salient evidence of the heterogeneity in rationales and motives underlying individuals' different privacy stances—consistent with work highlighting not just the diversity in privacy definitions [35], but also in privacy motives [31]. Third, it provides insight for policy makers regarding the extent to which claims of privacy “not being that important to me” may reflect an individual's simple disregard for privacy, or may have in fact more complex and nuanced interpretations.

2 BACKGROUND AND RELATED WORK

In this section we first review key theories of privacy behavior engagement (2.1), followed by the literature on privacy decision making (2.2) and the privacy paradox (2.3). We conclude with a discussion of existing works related to an inverse gap between privacy perspectives and behaviors (2.4).

2.1 Privacy Behavior Engagement

As this manuscript is concerned with the exploration of possible gaps between privacy perspectives and actual privacy choices, we start by considering two major theories of privacy behavior. Privacy Regulation Theory (PRT) [10] and Communication Privacy Management (CPM) [32] provide fundamental theoretical underpinnings for understanding engagement in behavior, and form the basis for our operationalization of privacy-regulating behaviors in the studies we conducted.

Altman's PRT [10] analyzes at a high level of abstraction engagement with behaviors that may increase or reduce one's privacy. Altman construes privacy as a selective control of access to the self or one's group. Altman looks at privacy behaviors as a dialectic process in which an individual dynamically defines boundaries—a *privacy regulation* process, not merely a process of hiding, secluding, or protecting personal information. These behaviors are common in everyday life, and often engaged in without much conscious deliberation—such as leaving a group of people to take a personal call on a cell phone, or closing the door of one's office to have a private conversation with another person. The theory accounts for privacy existing at different levels (self vs group) and for it being a temporally dynamic process. PRT gives high consideration to the role of context (such as with whom you would engage in the behavior, when, and under which conditions). Accordingly, PRT allows for the same person to make completely different decisions

at different times, if context changes and if they so desire. The theory also accounts for the gap, or mismatch, that may exist between one's desired privacy and one's actually achieved privacy.

CPM [32] also adopts the concept of boundaries to describe how individuals make decisions whether to share or conceal information. CPM's rule-based management system accounts for the characteristics that define a privacy rule, establishing that an act of disclosure will happen following one's cultural norms for privacy and openness, gender, context, motivation, and risk/benefit trade-offs.

While Petronio's CPM focuses on communication and sharing or concealing of information, Altman's privacy-regulating behaviors encompass a wide range of human actions. Privacy as a selective control of access to the self implies a negotiation, or regulation, of boundaries between the self and others that may take numerous forms, including informational, spatial, or bodily boundaries. Key to our analysis of privacy-seeking or privacy-regulating behaviors is the recognition of the dialectic nature of the process of boundary regulation: we exercise privacy by alternately opening and closing the self to others—for instance, when we choose to share a dark secret with a close friend, but conceal it from those we do not trust. This conceptualization of privacy goes beyond simplistic representations of privacy that construe it merely as a static condition of hiding or seclusion.

Common across Petronio's CPM and Altman's PRT is the recognition of the fundamental role of context in privacy behavior. As noted, PRT allows for the same person to make completely different privacy decisions at different times. Conceiving privacy as a dynamic process implies that the boundaries of self and others change all the time under the influence of multiple factors. Those boundaries change from person to person, and from culture to culture, which means that the specificity of not just what is private, but also what *privacy itself* is will necessarily change across individuals and scenarios. This definitional diversity or richness is well documented in the literature [35], and is key to our empirical analysis, as it allows for the possibility that, when individuals claim to care or not to care for "privacy," such claims are best interpreted as reflecting that individual's personal and unique conception of private and public spheres, which may or may not overlap with other individuals' conceptions.

2.2 Privacy Decision-Making

A key element of Altman's PRT that influences our empirical analysis is the recognition that gaps may exist between a person's desired degree of privacy and what she can actually achieve. Such gaps have been extensively analysed in the literature on privacy decision-making, and are key to understanding both the debate surrounding the so-called privacy paradox (see, below, Section 2.3) as well as our analysis of a potential, inverted mismatch between (dismissive) claims regarding privacy and (protective) privacy behaviors.

An extensive body of work on privacy decision-making has highlighted obstacles and hurdles individuals have to contend with to achieve desired levels of privacy. One such obstacle is information asymmetry [3], which arises when there is a knowledge difference between the parties involved in the decision-making process. Privacy policies are often so long and complex that they fail to

inform consumers who are naturally limited by bounded rationality [3]. One way that people can work around bounded rationality is by using rules of thumbs and heuristics when making their decisions [3, 12]. These in turn are susceptible to cognitive biases. In fact, consumers can be influenced toward decisions that do not align with their preferences by interface designs that lead their decision-making process in a particular direction through framing effects [3], suggestions [21], and status quo bias [12, 24] — a phenomenon recently popularized under the term "dark patterns" [2]. Other biases relate to systemic hurdles in decision-making. For example, hyperbolic discounting [1] can lead us to trade long-term risk (e.g., privacy) for short-term gain (e.g., more easily connecting with friends). These behavioral hurdles (as well as objective economic hurdles described in [6]) can make it harder for individuals to achieve the degree of privacy they aspire to, and hence provide possible explanations for the so-called privacy "paradox" as well as potential reasons why the inverse mismatch we study in this manuscript may arise.

2.3 The Privacy Paradox

The purported dichotomy between people's self-reported preferences towards privacy and allegedly privacy-negating behaviors is often referred to as the "privacy paradox" [29]. The privacy paradox has been fertile ground for research. A vast array of explanations have been proposed over the years. Early explanations focused on the fact that generic privacy attitudes, broad in nature, cannot be expected to predict actual privacy behaviors, which are contextual, specific, and nuanced [8]. Over time, as evidence of gaps and mismatches kept arising between specific mental states and specific behaviors also arose, other, non-mutually exclusive explanations kept being offered, derived from social theory, cognitive biases and heuristics in decision-making, decision-making under bounded rationality and information asymmetry, privacy calculus, and even quantum theory homomorphism [6, 23].

The privacy paradox has also been fertile ground for debate and controversy. Nearly twenty years of debate have failed to conclude whether it is a "myth" [36] or alive and well [13], as empirical researchers keep suggesting with new studies. We believe there are three reasons for the confusion, and each of them can tell us something useful for the research presented in this manuscript, by helping us identify pitfalls to avoid in our experimental design and analysis.

The first reason, as Acquisti, Brandimarte, and Lowenstein recently wrote [6], is merely lexical. The term "paradox" has different and subtly contrasting meanings. It can mean a self-contradictory statement that at first seems true, but also a statement that is *seemingly* (emphasis added) contradictory and yet is perhaps true (Merriam-Webster). The difference is important, and so is the emphasis on "seemingly." Scholars who interpret the term paradox under the first meaning look at possible explanations for gaps reported in empirical research between privacy mental states and actual behaviors, and because explanations do exist (many have been identified in the literature [5]), they conclude that there is no contradiction, and therefore no paradox (see, for instance, Solove (2020) [36])—even though, in fact, the evidence for the underlying gaps is robust. Whereas scholars who interpret the term paradox

under the second meaning look at the “seemingly” contradictory evidence of desires for privacy vs. privacy-denying actions, and call that a paradox (see, for instance, [13])—even though those gaps have by now well-researched explanations. We are agnostic towards the usage of the term “paradox.” In fact, we believe the focus on the *paradoxical* (or not paradoxical) nature of a gap between privacy perspectives and behaviors to be misguided: the debate around the privacy “paradox” being a myth or a reality arises principally from the lexicological ambiguity we just highlighted; ultimately, the paradoxical nature of a possible gap between claims and behaviors is in the eye of the beholder. Much more important (both to understand consumer privacy decision making and to inform policy) is the investigation of whether the gap itself does in fact exist and, if so, why. In this regard, much empirical evidence has suggested that mismatches between privacy-seeking perspectives and actual behaviors can and do in fact arise across multiple scenarios (which does not mean they will always arise: privacy behavior is contextual). For instance, specific attitudes toward app privacy did not predict app download behavior [11]; specific concerns over public self-disclosures did not match actual public self-disclosure behaviors [7]; specific expectations of privacy in social media did not match actual social media sharing behaviors [25]; even specific behavioral intentions were found not to match corresponding behaviors [29]. A vast array of factors has been used to explain those gaps, or mismatches. But the fact that misalignments between attitudes and behaviors can be explained does not mean that the misalignments themselves do not in fact exist [6]. Accordingly, in this manuscript, we are less interested in establishing whether a possible reverse gap may or may not be paradoxical, and we rather focus on whether it exists or not and, if it does, what could be its possible explanations.

Second, some of the confusion surrounding the privacy paradox is due to the fact that empirical research has searched for it across strikingly differing scenarios (with some naturally producing evidence for and evidence against the existence of gaps; contrast, for instance, the results in [7, 11, 25, 29] to those in [17, 39]), and under varying definitions (such as construing the gap as a mismatch between attitudes vs. intentions, or concerns vs. behaviors, or behavioral intentions vs. actual behaviors, and so forth). It stands to reason that mismatches and misalignments may occur under certain conditions and not others—sometimes, in the very same experiment [7]. In fact, some of the confusion specifically arises from framing the study of the relationship between privacy perspectives and behaviors in stark (and unreasonable) binary terms: *is the privacy paradox real, or is it an invention?* In reality, the relationship between privacy mental states and corresponding privacy behaviors is nuanced; it is as implausible to expect that privacy attitudes always predict behaviors, as to expect that they never do. In the studies presented in this manuscript, we account for these variations and nuances by taking an encompassing view of the potential gap between privacy perspectives (which may include attitudes, preferences, desires, concerns, and so forth) and privacy behaviors. Further, we explore both expressions of “general” privacy perspectives (such as “I have nothing to hide”) as well as perspectives on privacy in specific contexts, and we also take an encompassing view of privacy behaviors.

Third, as noted in the Introduction, some of the confusion surrounding the privacy paradox arises from the opposite implications for public policy that different observers have derived from it. From that confusion, we learn the need to be clear about the possible implications we personally draw from the evidence of a reverse gap between privacy perspective and privacy seeking behaviors. Centrally, the policy implication we derive from the evidence of a reverse gap is a caveat for policy makers to carefully vet whether claims of privacy not being any longer important truly reflect individuals’ underlying preferences for public policy, or may rather have more complex and nuanced explanations.

2.4 Work Related to a Reverse Privacy Paradox

Two manuscripts are relevant to our work. Adorjan and Ricciardelli used focus groups to study Canadian teenagers’ impression management strategies in social network sites, and uncovered evidence of a mismatch between expressing a “I have nothing to hide” perspective and engaging in privacy management on the sites. The authors concluded that the mismatch was the product of an adaptation to circumstances created by widespread use of social media [9]. Sun et al. investigated self-reported disclosure intentions and actual disclosures in IoT environments, and found evidence of a “reversed intention-behavior gap in IoT,” with participants disclosing less in the behavior condition than in the intention condition [37].

A few other manuscripts are tangentially related to the work we present here. In a qualitative study of cultural and generational influences on privacy concerns, Miltgen and Peyrat-Guillard found that younger individuals in their sample reported lower privacy concerns but greater protective behaviours [26], leading the authors to suggest that “additional research be conducted to further explore and unpack these relationships.” Ghaiumy et al. (2021) examined the privacy decision-making of older vs. younger adults and found older adults to make “more rationally calculated decisions than younger adults [...] negat[ing] the mainstream narrative that older adults are less privacy-conscious than younger adults” [20].¹

Differently from the above manuscripts, our studies are designed to focus on investigating the incidence as well as possible explanations for a reverse gap between several perspectives (preferences, attitudes, concerns, and so forth) and behaviors, across a large variety of settings, conditions, and actions. In so doing, we find that evidence of a reverse gap is actually common across our sample of participants; it appears across multiple scenarios; and it originates from various, distinct, factors.

3 WHY A MISMATCH BETWEEN PRIVACY PERSPECTIVES AND PROTECTIVE BEHAVIORS MAY ARISE

As noted in Section 1, the possibility of a mismatch between a person’s dismissive perspectives towards privacy and their privacy-regulating and privacy-protecting behaviors arises from a combination of anecdotes, common sense, and prior research. The previous sections, as well as the extant literature on the privacy paradox, suggest a tentative set of reasons why such a mismatch may arise,

¹ Wittes and Kohse (2017) write about a “privacy paradox II.” With that term, however, they refer to the paradoxical benefits of privacy threats [41], rather than the behavioral phenomenon we study here.

and why dismissive perspectives and privacy-seeking behaviors may in fact co-exist.

A first potential explanation arises from cases where an individual may engage in a behavior that appears to be privacy-protective, but the purpose behind the behavior has, in fact, nothing to do with privacy—even under a broad and encompassing Altmanian notion of boundary regulation. For example, upon seeing a participant clearing their browser history, a privacy researcher might assume that this was done to prevent others to see what the participant had viewed online. However, the participant may do this because they think it frees up space on their computer’s hard drive.

A second explanation arises from the fundamental difference between an Altmanian view of privacy as boundary regulation and individuals’ specific constructions of privacy in their everyday life. Altman’s view of privacy allows both the type of boundary and the specific “placement” of those boundaries to change from context to context, and from person to person—thus encompassing the richness and diversity of definitions of (from intimacy to seclusion; from autonomy to obscurity) and motives for privacy [31, 35]. Instead, as individuals, each of us construes privacy under our personal and idiosyncratic notions of private and public, and our own unique mental models of what privacy is (and is not). Those notions, and those models, often do not recognize or even admit for alternative ones—that is, for the notions and models that *other people* have of privacy. Thus, an individual will or will not interpret a statement as being associated with privacy depending on whether that statement matches her own distinctive conceptualization of privacy; this in turn will affect the extent to which that individual claims to be generally concerned about privacy. For example, when asked whether they are concerned about privacy in general, a participant who predominantly associates privacy with personal information, but does not care about the protection of her data, may reply that they are *not* concerned. And yet, the same participant may be highly concerned about their ability to maintain desired levels of seclusion from others. However, when asked about *those* concerns, that participant may respond that they have nothing to do with privacy—and may then reiterate their claim of not being concerned over privacy. And yet the next participant may exhibit the very opposite pattern of concerns, claims, and notions of privacy.

A third explanation arises from lessons learned in the privacy paradox literature, and focuses on the gap between general statements about privacy and specific instances of behaviors [8]. A participant may state that they are not concerned about privacy *generally*. That same participant may have concerns about specific scenarios and take protective actions to address them. For example, a participant might agree with the statement that they have “nothing to hide.” However, upon further questioning, the participant might add that ever since a nosy neighbor asked them about a conversation overheard in their backyard, they have taken to going inside when the neighbor is outside to avoid the neighbor’s meddling. Thus, this participant would acknowledge engaging in a specific instance of a privacy protective behavior even though they generally do not care that much about privacy.

A fourth possible explanation accounts for the fact individuals may not express concerns about privacy in a particular situation if they already engage in protective behaviors that they believe address their privacy concerns. This explanation focuses on the

feedback loop present in the Theory of Reasoned Action [19]. For example, a participant who uses web browser privacy plug-ins that block third-party cookies and other trackers may state that they are not concerned about advertisers tracking them online. In this case, the behavior is the reason why the concern is not present.

As is the case for the privacy paradox literature, there may be other, and certainly not mutually exclusive, explanations for possible mismatches between dismissive privacy stances and actual engagement in privacy-regulating and privacy-protective behaviors. The analysis we present below aims at exploring evidence of such mismatches and which of the many possible explanations may be empirically validated.

4 BEHAVIOR ELICITATION STUDY

We conducted a series of online studies between February and July 2021. We first conducted a Behavior Elicitation Study (discussed in this Section), to collect a list of privacy-regulating and privacy-protective behaviors from non-experts. Next, in Study 1 (Section 5), we used the collected list of behaviors to explore gaps between privacy perspectives and self-reported behaviors. Finally, in Study 2 (Section 6), we explored potential explanations for observed gaps. The combination of studies allows us to investigate the incidence of reverse perspective/behavior gaps (RQ1), and their possible explanations (RQ2). All studies were approved by Carnegie Mellon University’s Institutional Review Board. All participants were US residents recruited on Prolific Academic and were at least 18 years old. Participants across all studies provided informed consent.

4.1 Methodology

The behavior elicitation study aimed at collecting self-reported examples of privacy-protective behaviors from non-experts, both to generate a broad and encompassing set of behaviors that would embody and reflect the richness and diversity of individuals’ notions of privacy, and to identify behaviors beyond those traditionally used in the literature, typically generated by experts. We recruited participants to complete “a survey about your everyday behaviors in a number of different contexts.” Participants took an average of 20 minutes to complete the survey and were paid US\$5.

The behavior elicitation study used a 3x2 between-subjects survey design. First, we asked participants to identify situations in which they felt the need to protect their privacy; then we asked them which steps, if any, they took as a consequence of the situation they mentioned. Participants were randomly assigned to one of six conditions, with the prompt to the first question manipulating the *context* of the question (x3: offline, online, undefined) and the *prominence* of the term privacy (x2: direct, indirect). In addition, within subjects, we asked participants to consider multiple *actors* relative to which the participants may have wanted to regulate their private boundaries, including others, the government, companies, malicious actors, and strangers. This design casts a broad net for capturing potentially privacy-relevant behaviors in a wide range of scenarios and contexts, while accounting for potential priming concerns caused by the use of the term privacy.

Table 1 presents the specific wordings used in each framing, which were refined through iterative tests conducted on Prolific. As an example, here is the wording for a question involving a

Variable	Framing	Wording
Prominence	Direct	“to maintain, manage, or protect your privacy.”
	Indirect	“to manage or control the degree to which others can observe and learn about you?”
Context	Online	“when [you are] using the Internet”
	Offline	“when [you are] not using the Internet”
	Undefined	“in your everyday life”
Actors	Open ended	“from others”
	Government	“from government surveillance technology”
	Acquaintances	“from people you know”
	Companies	“from companies that collect your data”
	Malicious actors	“from malicious actors, like hackers or other ill-intended individuals”
	Other people	“from strangers”

Table 1: Description of the different framings and wordings used (Behavior Elicitation Study).

government actor asked to a participant assigned to see an indirect prompt with an online context: *Please describe situations, when using the Internet, in which you feel (or felt) the need to manage or control the degree to which government surveillance technology can observe and learn about you.* This was followed by: *For each situation above, what steps do (or did) you take as a consequence of that situation?*

Participants were offered two separate input boxes per prompt so that they had the opportunity to provide two situations and steps related to each. However, they were only required to provide one example per question to continue with the survey.

We used a qualitative coding method, described below, to iteratively develop clusters of behaviors from participant responses before analyzing them.

4.2 Results

The behavior elicitation study was completed by 60 US Prolific participants. Of them, 28 self-reported as male, 31 as female, and one preferred not to answer. The self-reported ages ranged from 19 to 56 years, with the average age of 32 years. As each participant provided up to two responses to scenarios considering six different actors, we collected 617 pairs of privacy-concerning situations and the steps taken by the participant(s) to address them. We removed responses associated with security or safety (55), and responses that did not follow the request to provide a privacy-protective behavior taken as a consequence of the situation presented (52). We also separated responses in which participants presented multiple behaviors for the same privacy-concerning situation. As a consequence of this cleanup process, we analyzed a total of 597 valid responses to extract privacy-protective behaviors.

From the 597 valid responses, we iteratively developed a codebook of categories of privacy-protective behaviors, systematically assigned codes to each response, clustered responses based on these codes, and identified clusters that could be merged under a single,

more abstract code. We repeated this process until the set of behavior codes did not contain significant overlap.²

The outcome of this process is a list of 77 behaviors: 73 extracted from the study, and four added by the research team that represent modifications, in a different context, of behaviors provided in participants’ responses. The final list of behaviors (Table 6 in the Appendix) is a corpus that we expect to be useful in future studies.

5 STUDY 1: BEHAVIOR ENGAGEMENT AND PRIVACY PERSPECTIVES

Study 1 explored gaps between privacy perspectives and self-reported behaviors using a within-subjects survey experiment. The study design tackled RQ1: is there general evidence of gaps or mismatches between (dismissive) privacy claims and (protective) privacy behaviors? In addition, it served to identify participants for Study 2, in which we explored the reasons behind mismatches.

5.1 Methodology

Study 1 used a within-subject, two-part survey design. We recruited participants for “a survey about your behaviors in a number of different contexts” and compensated each with 75 cents in Part A and \$1 in Part B. Participants took an average of 3 minutes to complete Part A and 4 minutes to complete Part B. In Part A, we assigned participants to a subset of privacy behaviors randomly selected among the ones collected during the Behavior Elicitation Study, and then collected data on participants’ engagement with those behaviors. In part B, three days later, we collected participants’ privacy perspectives.

Part A: Behavior Engagement. In Part A, each participant was randomly assigned to a selection of ten protective behaviors drawn from the final list of 77 behaviors obtained from the Behavior Elicitation Study. For those ten behaviors, participants provided their engagement on a 6-point frequency scale: Never, Rarely, Sometimes, Frequently, Most of the time, Every time the opportunity presents itself. They were also offered the options “I don’t know” and “Does not apply,” to separate participants who had never engaged in the behavior by choice from those that had not due to circumstance.

Participants were also asked to explain their underlying motivation for engaging in the behavior through the following prompt: “If you had to explain why YOU do each of the behaviors below, what would YOUR primary reason be?” Participants were provided with the following options and asked to pick one: safety, social norms, convenience, practical reasons, privacy, security, and another reason. For those who selected “another reason,” we asked them to describe it.

Part B: Privacy Perspectives. Three days after Part A was completed, we invited Part A participants to return for Part B by sending a message through the Prolific platform about “a similar study” called “Study about attitudes” that would be available in their Prolific interface. Part B was conducted separately to reduce the potential for bias in responses had participants been asked to answer both about their perspectives and behaviors in the same survey.

²The first author developed the codebook, which was then iteratively refined with all authors following an open-coding process for open-ended responses. The first author also coded the responses in both Study 1 and Study 2.

	Gender (%)				Age				Behaviors engaged with
	Female	Male	Non-binary	No response	Mean	Median	Min	Max	Mean
Study 1A	54	42	3	1	31.6	29	18	77	7.7
Study 1B	59	36	5		31.7	29	18	64	NA
Study 2	63	33	4		31.5	28	18	64	7.9

Table 2: Participant demographics (Study 1 and Study 2)

In Part B, we presented participants with 25 general statements surrounding their attitudes, concerns, and preferences related to privacy. The vast majority of the statements were generated from existing privacy scales [34]. In addition, as existing validated privacy scales do not cleanly differentiate between attitudes, concerns, and preferences, we generated some additional statements based on recent work on how to best construct statements to capture privacy perspectives [15]. This approach is appropriate for our purpose of offering a range of general attitude, concern, and preference statements to compare with self-reported behaviors. We did not use the statements collectively as a scale. The general statements can be found in Table 3.

In addition, we presented participants with ten statements of preference and ten statements of concern that directly matched the ten behaviors that we asked them about in Part A (preference and concern statements are shown in Table 6 in the Appendix.) This led to 45 perspective statements presented in a random order to participants. They reported their agreement with each statement in 5-point Likert items, ranging from “Strongly disagree” to “Strongly agree.”

Participants who presented a *dismissive* perspective,³ as captured by any of the 25 general statements, were provisionally classified as “someone who might not be particularly concerned or care about privacy.” We asked these participants for the degree of agreement with the characterization. If they did not recognize themselves in that characterization, we asked them to explain why.

5.2 Results

Part A of Study 1 was completed by 255 participants. Due to attrition, Part B was completed by 73 participants. As shown in Table 2, our sample skewed slightly non-male and young.

5.2.1 Part A: Behavior Engagement. Participants reported having engaged, on average, in 7.7 behaviors out of the 10 presented to them at least once. Aggregating all participants and behaviors, only 21.5% of the responses reported no previous engagement with the behaviors, and only 2% fell in the category of no previous engagement due to lack of awareness or applicability. This meant that the vast majority of participants reported engaging in a *majority* of the behaviors that we collected in our behavior elicitation study.

We observed that aggregating across all behaviors in which participants reported some engagement, 41% of participants reported

³By *dismissive* perspective, we refer to a perspective with negative affective valence towards privacy: For statements that would indicate a positive affective valence perspective towards privacy, the participant strongly or somewhat disagreed; for statements that would indicate a negative affective valence perspective towards privacy, the participant strongly or somewhat agreed.

Statement	Agreement (%)		
	D	N	A
I don't think that privacy is important for society	95	4	1
I don't think that privacy is important to me	92	7	1
Only people who have something to hide need privacy	85	7	8
I think that others worry too much about privacy	75	10	15
My life is an open book	60	16	23
I don't mind that others know what I'm doing	59	21	21
I want to be able to keep at least some aspects of my life private	0	1	99
I want to have the ability to be left alone	0	4	96
I do not want others accessing my accounts or devices without my permission	0	4	96
I do not want to be surveilled or tracked	0	4	96
I want to be able to control what others learn about me	1	8	90
I want to have control over how my data is collected or used	3	7	90
I worry about others getting access to my accounts, devices, or information without my permission	3	8	89
I worry about companies selling or sharing my personal information without my knowledge	5	8	86
I worry about my personal information being leaked	3	18	79
I worry about others observing me without my knowledge or permission	8	15	77
I worry about being tracked online or offline	7	18	75
I worry about the fact that too much of my data is being collected	14	11	75
I worry about not having privacy from the government	15	10	75
I worry about identity theft	7	19	74
I'm concerned that we, as a society, will lose our privacy	8	19	73
I worry about not being able to have privacy anymore	14	14	73
I worry about losing my ability to be left alone by others	18	18	64
I'm uneasy about the current amount of privacy I have	25	19	56
I worry about some of the people I know learning about certain activities or opinions I have	26	18	56

Table 3: Distribution of general statements and general perspectives (Study 1, Part B). For the first six statements, agreement indicates a negative privacy attitude. Agreement with the next six statements represents a positive preference for privacy. Agreement with the remaining 13 statements indicates a privacy concern. D stands for disagreement, N for neutral, and A for agreement.

low engagement, with 16% of participants reporting engaging with each behavior “Rarely,” and 25% “Sometimes.” On the other hand, 36% of participants reported high engagement, with 12% reporting engaging with each behavior frequently, 11% most of the time, and 13% every time the opportunity presents itself. (The remaining 23% included “Never,” “I don’t know,” and “Does not apply” responses)

When we asked participants to categorize their motivation to engage in each behavior their most commonly selected motivation overall was “Privacy” (30%) and “security” (17%). Table 4 shows the engagement frequency for the behaviors where more than 50% of participants selected privacy as the main motivator for engaging in the behavior. That noted, even though the behaviors were collected through a survey designed to elicit privacy-protective behaviors, only 30% of responses in Part A categorized the motivation to engage in those behaviors as being *primarily* associated with privacy. This finding may indicate both that different people engage in the same behaviors for different reasons, and that—consistent with the possibility, considered in Section 3, of heterogeneous and idiosyncratic notions of privacy across individuals—they may construe the very term privacy differently. We consider these possibilities further in Section 6.2.

The counts of participants’ self-reported frequency of behavior engagement and motivation are available in Table 7 in the Appendix so that they may be used by researchers in future studies.

5.2.2 Part B: Privacy Perspectives. Out of the 255 participants, 111 responded to the invitation for Part B. We piloted Part B with a batch of 37 participants to ensure task comprehension, response quality, and survey duration. We then made some changes to the wording of the survey as a result of this analysis, and discarded those pilot responses. We also discarded one participant whom we identified as providing low-quality responses,⁴ leading to a total of 73 participants completing Part B. We paid participants US\$1 for Part B, which took on average around 4 minutes to complete.

Participants who returned for Part B had previously reported engaging in 7.9 out of the 10 behaviors at least once. They reported negative/dismissive views towards privacy, on average, on 4.3 out of the 45 statements meant to capture their privacy perspectives.

Based on the 25 general statements related to our participants’ attitudes, concerns, and preferences about privacy, we identified a majority of our participants as expressing positive views on privacy, with at least 75% of participants taking a privacy-positive view on 17 of the 25 statements (see Table 3). Nevertheless, some of the statements showed lower rates of privacy-positive responses. In particular, only about 60% of participants disagreed with the statements “My life is an open book” and “I don’t mind that others know what I’m doing.” Furthermore, participants less frequently agreed with the statements of worry.

We observe similar patterns for the statements about preferences and concerns associated with specific behaviors. Only three specific preference statements did not have privacy-positive responses from a majority of participants: “I do not want others to find out about my activities in certain places” (50% disagreed), “I do not want to be tagged on my friends’ posts” (50% disagreed), and “I do not want to be recorded in security camera footage while in public” (63% disagreed). For the specific concern statements, there were five

⁴Their answers did not actually address the posited question.

statements where a majority of participants did not have privacy-positive responses: “I worry about government surveillance” (50% disagreed), “I worry about my friends tagging me on their posts” (50% disagreed), “I worry about not being able to be left alone” (60% disagreed), “I worry about security cameras in public areas” (75% disagreed), and “I worry about my privacy when shopping in person” (82% disagreed).

5.2.3 Comparison of Behaviors (Part A) and Privacy Perspectives (Part B). By comparing their Part A and Part B responses, we aimed at exploring a potential inverse gap, or mismatch between privacy claims and privacy behaviors in our sample.

We define four categories of potential mismatch:

- **General attitude mismatch:** the participant reported engaging in at least one protective behavior but agreed (somewhat or strongly) with at least one of the *general attitude* statements expressing a negative attitude towards privacy.
- **General preference and concern mismatch:** the participant reported engaging in at least one protective behavior but disagreed (somewhat or strongly) with at least one of the *general statements regarding privacy preference and/or concern*.
- **Specific preference mismatch:** the participant reported engaging with a protective behavior but disagreed (somewhat or strongly) with the *preference statement associated with that behavior*.
- **Specific concern mismatch:** the participant reported engaging with a protective behavior but disagreed (somewhat or strongly) with the *concern statement associated with that behavior*.

We observed at least one type of mismatch in over 70% of the sample who completed Part B (52 participants out of 73). Overall, 28 participants exhibited a general attitude mismatch, 35 a general preference and concern mismatch, 33 a specific concern mismatch, and 19 a specific preference mismatch. As these numbers indicate, some participants displayed multiple types of the mismatch.

After filtering out the behaviors that the participant did not see as privacy-related in Part A (that is, a behavior for which the participant did not select “privacy” as the *main* motivator), we see a small reduction in the incidence of mismatches as a whole (from 71% to 67% of participants), and a more significant reduction in the incidence of specific concern mismatches (from 45% to 21% of participants). Table 5 summarizes the distribution of different types of mismatches generally, and with privacy as the main motivation behind the behavior.

In summary, Study 1 indicates that engagement in behaviors that can be construed as privacy-seeking is very common among our participants (Part A), and that mismatches between privacy-seeking behaviors and negative, dismissive privacy claims are common as well, although those mismatches take different forms (Part B).

6 STUDY 2: EXPLANATIONS

Whereas Study 1 was a quantitative study aimed at exploring whether gaps between dismissive claims regarding privacy and actual engagement in protective behaviors exist, Study 2 was a qualitative study aimed at exploring potential explanations for those gaps when they arise.

Privacy as motivator (%)	Behavior	Engagement (%)		
		None	Low	High
71	Made sure I was alone when using the internet for private matters like gift shopping, private emails, porn, etc	13	29	58
69	Gave the least amount of information when answering a personal question.	0	69	31
65	Reprimanded someone for behavior that invaded my privacy, such as entering my private room or looking through my stuff.	12	74	15
62	Took steps so that others couldn't see what I was doing, for example, minimized my screen, sat with my back to a wall, shielded my phone or documents, used a screen protector, etc.	8	59	32
61	Changed settings or blocked people on social media to only allow certain people to view my information (posts, photos, etc.).	6	58	36
58	Took steps so that services, apps, or websites couldn't access my location without my knowledge, for example, turned my location off, or kept it off, turned my phone off, etc.	3	44	53
56	Used tools to increase my privacy like private browsing or incognito mode, blocking third-party cookies on my browser, encrypted emails, VPN, adblockers, etc.	3	47	50
51	Attempted to hide my internet footprint, for example, by deleting information about me.	17	74	9
51	Changed the status indicator on an app or website to not show I was online.	24	43	32
51	Spoke more quietly, moved away, or wrote something down instead of saying it to avoid others from overhearing a sensitive conversation.	19	51	30

Table 4: Engagement frequency for the behaviors where more than 50% of participants selected privacy as the main motivator (Study 1, Part A). Engagement categories were combined: None (Never, I don't know, Does not apply), Low (Rarely, Sometimes), High (Frequently, Most of the time, Every time the opportunity presented itself).

Type of mismatch	Motivations (%)	
	All	Privacy
<i>At least one type of mismatch</i>		
All participants	71	67
Participants who agreed with assessment	27	23
<i>Specific types of mismatch</i>		
General attitude	38	–
General preference and concern	48	–
Specific concern	45	21
Specific preference	26	16

Table 5: Distribution of the different types of mismatch for all behaviors, independent of motivation, and only for behaviors originally classified as privacy (Study 1).

6.1 Methodology

Study 2 focused on the subset of Study 1 participants who had reported engaging in at least one of the 10 behaviors from the first survey but had presented a negative, dismissive perspective on privacy, either in the general statements or in terms of a mismatched perspective for the preference and concern statements matched to those engaged behaviors.

The subset of fifty-two participants (70%) from Study 1 who exhibited a general or a specific mismatch were invited to proceed to a “bonus” study for an additional \$2 compensation; 49 agreed and were immediately presented with the prompts for Study 2. The survey took on average about 12 minutes to complete.

For each behavior that they had engaged in, participants were asked to explain *why* they thought a mismatch had been observed through a combination of close-ended and open-ended survey questions. For **general mismatches**, participants were presented with multiple-choice explanation options, as well as an open-ended text field. In addition, participants were asked whether they agreed or disagreed with a characterization of them as “someone who might not be particularly concerned or care about privacy” based on their Part 2 responses. Participants who disagreed with the characterization were asked to answer the general mismatch questions about potential factors beyond that disagreement. For **specific preference and concern mismatches**, participants were given an open-ended text field to provide their responses.⁵

We followed an open-coding process for open-ended responses: the first author analyzed each response in relation to the behavior and statement they referred to and assigned it a high-level descriptive code. These codes were refined and responses were grouped based on their similarity.

6.2 Results

We focus on the qualitative results from open coding, in keeping with our study objective of identifying potential explanations of mismatches. We present numbers to offer a summary description

⁵We offered multiple-choice options for general mismatches and only open-ended questions for specific mismatches to decrease participants fatigue and avoid the risk of reducing the overall quality of responses (answering every possible mismatch may significantly extend the time necessary to complete the survey and increase participants' cognitive costs). Also, we expected specific mismatches to originate from context-specific, nuanced factors that made the use of open-ended answers more compelling than for general mismatches.

of our data set, but given the small and self-selecting nature of our sample, we refrain from drawing quantitative conclusions.

6.2.1 General mismatches. A general mismatch arises when a participant claims to have engaged in *at least* one behavior but expresses *at least* one negative general privacy perspective when answering the general attitude statements or general preference and concern statements.⁶ Forty-six participants had exhibited a general mismatch in Study 1. In Study 2, those participants were asked to discuss why, in their opinion, a mismatch between reported behaviors and perspectives had been reported. Participants were given five options as well as an open-ended field and asked to check all answers that applied. The most common reason to explain why the mismatch occurred was that engaging in the behavior reduced their level of privacy concern (43% of general mismatches). The second most common reason was that there are sometimes situations where engaging in the protective behavior is warranted (40%), followed by having engaged in the behavior for a reason other than privacy (22%), having answered the statements about privacy perspectives thinking about privacy in a different way (15%), and “Other” (14%). The vast majority of the open-ended explanations they provided under “Other” matched having engaged in the behavior for a reason other than their notion of privacy.

Roughly a third of participants who had engaged in at least one privacy behavior nevertheless agreed with a characterization of them as someone who did not care or was not concerned about privacy. Just over two thirds (32 out of 46) disagreed. Participants who disagreed expressed negative perspectives on an average of 2.3 out of the 25 general statements while those who agreed expressed negative perspectives on an average of 6.6 statements. After excluding participants who disagreed with the characterization, roughly one in four participants still exhibited evidence of general mismatch (27% of participants considering all motivations; 23% considering only behaviors where the participant listed privacy as the main motivation for engagement).

The explanations provided were frequently based on caring about privacy *in some forms* but not others, and not being *overly* concerned about privacy. As one participant put it:

I do care about privacy and security, but I don't obsess over it. I make sure I have strong passwords especially for websites that may use my credit card information. ...I don't share personal/identifying information like addresses, SSN, last name, etc. on the internet unless it's a trusted website. ...All that said, I'm not constantly worried about my privacy being compromised or paranoid about being tracked and surveilled.

6.2.2 Specific mismatches. Specific mismatches arise from engaging in privacy-seeking behaviors but expressing dismissive perspectives (self-reported attitudes, concerns, or preferences) associated with those behaviors. In Study 1, 33 out of 73 participants had exhibited a specific concern mismatch and 19 a specific preference mismatch. When asked to explain why we observed those mismatches, participants proposed numerous explanations.

The most frequent response involved having engaged in the behavior for a reason unrelated to the *participant's* notion of privacy. For instances of mismatches between behaviors and specific preferences, these responses accounted for about a quarter of the explanations; for specific concerns, they accounted for about half of the explanations. Motivators that were frequently reported by participants were convenience, practical reasons, mental health, or notification and spam reduction. In some cases, participants who argued they had engaged in the behavior for reasons unrelated to privacy offered motivations that indeed appear unrelated to privacy. For instance, when explaining the gap observed over the usage of a secondary email for some services, one participant stated “I don't worry about sharing my email much, I just like to keep things moderately organized.” In other cases, however, non-privacy motivations offered by participants seemed consistent with encompassing theories of privacy such as Altman's Privacy Regulation Theory [10]. For instance, one participant explained that they did not agree with the statement “I worry about discussing sensitive topics” by stating that “I don't care if people know where I stand on sensitive subjects.” But then added “I just don't feel like talking about certain subjects with people who have vastly different opinions than mine,” appearing to be willing to regulate their level of closeness and openness to others depending on context. We found similar occurrences of participants ruling out privacy as the main motivator while at the same time presenting motivations that would fall under privacy theories that do not limit the notion of privacy to mere protection of personal data. For example, one participant stated “I just want to do things on my own terms and at my own pace. It's not like I have anything to hide, but I just don't want to answer to anyone else as I enjoy some alone time.” This desire to be left alone and achieve solitude was echoed by multiple participants who stated that their motivation for engaging in the behavior was *not* privacy—while, of course, a right to be left alone has long been associated with privacy in the literature [40].

In some other cases, differences between a privacy motivation and the underlying non-privacy motivation for the behavior offered by the participants were complex to disentangle, given that the wordings used to describe some behaviors to the participants were explicit and specific about its underlying privacy connotations. For example, one participant had reported having chosen “to shop online instead of in-person due to *privacy concerns*,” but later claimed not to agree with the statement: “I worry about my privacy when shopping in person.” That participant explained the gap by providing an answer focused on convenience, highlighting that, “Convenience and safety during a pandemic are more important.” It may be the case that this person had read the statement in Part A quickly and not noticed the “due to privacy concerns” text, or it may be the case that the participant reconciled the potential mismatch by ranking privacy concerns on a varying scale, relative to other concerns, and based on circumstances.

Overall, these observations highlight how the remarkable heterogeneity of privacy notions, constructions, and definitions commonly observed among scholars (to the point of privacy being deemed a “concept in disarray” [35]) is not exclusive to experts: individuals may hold highly personal and diverse views of privacy or of what privacy is, and is not. Those views, notions, and *conceptual boundaries* of privacy differ both from subject to subject, and between

⁶As we consider this an exploratory study, we looked for evidence of *potential* mismatches—hence we picked a low threshold for the definition of mismatch.

subjects and experts. Ultimately—as discussed in Section 3— notions of privacy may be so exclusive to each individual that a participant may engage in behaviors for motivations that other individuals construe in fact as privacy-related, but the participant herself does not consider pertaining to privacy. Such a complex but rich milieu of privacy notions should be accounted for when interpreting results of privacy surveys, and when policy-makers attempt to understand the meaning of claims of not caring for privacy.

In other instances, participants concurred with the privacy motivation but had concerns that did not align with the statements presented to them. One participant stated that they had asked friends not to tag them in photos or posts in social media, not because they were concerned about the *act* of being tagged, but because of their *reaction* to the notifications: “I don’t like getting texts that say I’ve been tagged in something. I tend to overreact and rush to the website to see what’s been said about me.” At times, this misalignment between our presented concern and participants’ concerns was due to the statements’ specificity. For example, one participant reported not worrying “about others observing me while I’m at home” explaining that “It is more of a concern of someone coming into my home rather than peeking in the windows.” And yet this concern, again, appears consistent with an Altmanian view of privacy as a process of boundary regulation — in this case, applied to a person’s personal space.

Another common explanation for mismatches provided by participants was that their concerns and preferences were not as pervasive as the associated statements implied — they *were* concerned, yes, but only in *some* cases. For example, a participant who reported having “avoided using public or unsecured WiFi” but reported not being worried “about using public or unsecured WiFi” explained:

I’m usually fine with public WiFi but not for things like banking because although the chance of someone taking advantage of me is small, I’m not willing to risk all of my money even a little bit.

Another, who had reported *not* being worried about others being able to see them online, but had also reported having changed their status to not show that they were online, stated:

Most times I do not worry about what friends see me online. But there have been times where I did not want certain friends to see me online because they would be more likely to contact me and I was not in the mood to conversation.

These context-dependent worries and preferences varied along the lines of what information was involved (“I don’t mind strangers knowing any nonharmful information, I draw the line at telling people where I live for no reason.”); who was involved (“I do not want to share my private networks with strangers, but I’m not private with it for people I know.”); participant location (“I only do this in workplace environments to maintain good relationships with colleagues.”); and, why they found themselves in the situation associated with the behavior (“I normally want to be seen on security cameras because it would provide a good alibi if I’m falsely accused of a crime, but sometimes I don’t want to be seen.”). These comments appear consistent with the widely accepted notion in privacy scholarship that privacy preferences must be understood within their context [28]—but so also, it turns out, should claims of “not

caring” or “not being worried” about privacy. As one participant put it:

This behavior and attitude are both context-dependent. A stranger at a charity auction, for example, is someone I would treat differently from a person on the street in a bad neighborhood. My attitude reflects my hopes, while my behavior reflects some harsh realities.

Other explanations for a lack of worry related to our participants having taken steps to not expose themselves or to protect themselves. For example, one participant explained that they did not worry about others seeing their exposed body, even though they reported having covered themselves to avoid that because it depended on the situation in which others saw their body and that “there are a few ways that I do care and are careful.” Another, when faced with the same gap, said that “I don’t expose myself online, so I don’t worry about that. I do cover myself though when others are present.” One participant explained:

I don’t worry about people seeing my information because [of] all the settings I have in place. I don’t want them having full access to my profiles, but I don’t worry about it.

This participant highlighted another explanation that was not always openly expressed: that the source of worry was not significant, frequent, or disruptive enough to warrant becoming a concern. When explaining the lack of worry related to the trade of personal information for services like newsletters or apps, one participant said: “Newsletters and apps feel less invasive to me than other services, even though that’s likely untrue.” Others, when explaining their lack of worry in the context of unwanted calls explained, “It is not a big deal if I can confirm the call is real” and “It is not significantly disruptive to my daily routine to worry about.” In fact, for approximately 67% of instances of concern-based mismatches, participants reported disagreeing with the expression of concern but agreeing with the expression of preference for the same behavior. If we include neutral responses, this percentage increases to 83%. In addition, one participant, when explaining why they disagreed with “I worry about being spied on online” even though they stated having taken steps so that the government or other entities could not spy on them explained that:

Though I do not have much to hide, I do not feel the government should have access to this much data about my personal choices; it’s simply not necessary – It’s a stand based on my values, not my concerns.

Of note, a participant who stated they had previously avoided using facial recognition as an identification method but disagreed that they did “not want to use facial recognition technology,” explained that they did not necessarily see the technology as bad, but would try to avoid it until they could learn more. This suggests that the participant might have seen the statement as value-laden (implying the technology was “bad”) or they might consider their avoidance of the technology as a temporary measure while they learn more about it rather than a decision not to use it.

Lastly, some infrequent explanations were associated with the study design. For example, we had intended the behavior “Secured or shredded sensitive documents” and associated concern statement

to be related to physical documents (and the analogous behavior “Encrypted or password protected sensitive files” to be associated with digital documents). However, one participant stated that they disagreed with the sensitive documents concern statement because they were thinking about digital documents, stating that “I only worry about physical copies of documents.”

7 LIMITATIONS

We consider this set of studies an exploratory analysis of the phenomenon of a reverse privacy paradox. Before we discuss our results further, we want to highlight several limitations that affect the generalizability of our results. First, our sample of participants (US Prolific users) is not representative, and is US-centric. Some prior work has suggested that online crowdsourcing samples may be *more* privacy-sensitive than the rest of the population [22]. This may affect our results in various different ways: a *higher* frequency (relative to a nationally representative sample) of individuals reporting low privacy concerns because, in fact, they had already taken steps to mitigate their privacy risks; as well as a *lower* frequency of individuals actually reporting low privacy concerns. That noted, Prolific samples are commonly used in usable privacy research and have been found to be fairly representative for surveys about user perceptions and experiences [38]. Second, the estimates of the frequencies of detection vary depending on whether we adopt looser or stricter definitions of it (see Section 8). Our *definition* of general mismatches could be considered loose (as we considered it to occur whenever a participant engaged in at least one behavior but expressed *at least* one negative general privacy perspective, out of the many presented to them); whereas our count of general or specific mismatches may be considered overly strict: because we asked each participant about only 10 randomly selected behaviors, we are likely missing instances of mismatches involving behaviors that a particular participant may have engaged in but were not among the 10 presented to that participant; in addition, some participants may engage in privacy-protective behaviors not on our list. Finally, due to attrition, we experienced a significant drop out rate between Part A and Part B of Study 1, which also led to a relative small sample for Study 2. The small and self-selecting nature of this sample reduces our ability to generalize about the frequency of various explanations discussed. Future work should aim at expanding our exploratory analysis with larger and more diverse samples sizes.

8 DISCUSSION

Using quantitative and qualitative methods, we performed an exploratory investigation into a phenomenon where people present negative perspectives on privacy (such as not caring or not being concerned), but report engaging in privacy-protective behaviors. As noted in the Introduction and in Section 2, we sidestepped the issue of whether this reverse gap is ultimately paradoxical; instead, we focused on an empirical investigation of its existence and its possible explanations.

We found, first, that our sample of participants reported having a high engagement with the potentially privacy-protective behaviors presented to them. On average, our participants had engaged at least once with approximately eight of ten behaviors presented.

Second, by adopting an encompassing approach to identifying the phenomenon—namely, by considering any pair of negative perspectives and engaged protective behavior as an example of mismatch—we identified it as present, in either its generic or specific form, in the responses of a large share of participants. That share remains high also after conservatively filtering out behaviors where the participant had not initially selected privacy as a motivation for engaging in the behavior. After further excluding participants who disagreed with a characterization of their *general* views towards privacy as negative or dismissive, those percentages decreases to about one out of four participants. That noted, considering the limitations we have outlined in the previous section and the exploratory objectives of this study, we urge the reader to focus more on the broader implications of these results than on its exact frequency: privacy-seeking behaviors are quite commons among the respondents; some form of inverse privacy perspective/behavior gap does seem to occur in a sizeable fraction of respondents; and that gap has diverse and complex explanations that vary from respondent to respondent.

Third, we found that the majority of participants have a low percentage of mismatches, indicating that this phenomenon happens in specific instances for each individual, as opposed to being a general trait of the participant.

Fourth, common themes emerged from the analysis of possible explanations of the mismatches. Participants frequently explained the misalignment between specific preference and concern statements and their behaviors by highlighting that they engaged in privacy-protective behaviors in some specific situations even if they were not concerned about that entire category of potential situations. This reasoning is consistent with one of our potential explanations about the observed gap (Section 3): *participants were only concerned about specific instances of broader privacy concerning situations*. This observation also confirms the challenges associated with collecting people’s privacy perspectives such that we are true to their nuances. Future work on privacy surveys may benefit from carefully creating and balancing statements such that they are neither too generic or all-encompassing, where people might feel that their nuanced view on privacy is not being captured, nor too specific to the point where those answering the survey will have difficulty seeing the difference between statements. Furthermore, researchers should remain vigilant when interpreting results that include participants’ privacy perspectives, as participants might have responded differently if subtle variations had been included when capturing participants’ responses.

In addition, we found evidence that some participants engaged in a behavior for privacy-protective reasons while others engaged in the same behavior for reasons they did not consider related to privacy. In some cases, the reasons presented for engaging in the behavior were quite clearly not related to privacy even in an encompassing interpretation of the concept (for instance, reasons of convenience). In other cases, the reasons presented highlighted the rich and complex nature of individuals’ notions of privacy and are consistent with another proposed explanation for a mismatch: *people having varying, idiosyncratic, and even contrasting interpretations of privacy*. By analyzing participants’ open-ended responses to the question why the mismatch was observed, we found evidence of different interpretations of privacy among participants. Cases

in which participants reported engaging in a behavior for a reason unrelated to privacy, and yet provided a motivation that fits under privacy theories confirm the notion that individuals hold differing views of privacy. Thus, researchers need to be careful to ensure that they account for this variation in “privacy lens.”

We also observed some participants change their view of their underlying motivation for engaging in the behavior when explaining the mismatch. This highlights the challenges of interpreting individuals’ motivations for undertaking potentially privacy-preserving behaviors. In fact, many participants stated that they had engaged with a behavior (which included, in the description of it presented to participants, an explicit privacy motivation), only to later state that they did not see the behavior as privacy-motivated. These instances may highlight some participants’ attempts to reconcile, *ex post*, mismatches by ranking privacy concerns against other situational concerns; or may reflect a broader potential limitation of collecting this type of data via surveys—inattentiveness. Although, in our study, only one participant stated that they must have picked the wrong choice as an explanation for the observed gap, surveys are prone to “gut feeling” responses and, when asked to explain a mismatch, participants might reconsider that initial reaction. The latter was studied by Phelan et al. (2016) as an explanation for the privacy paradox [33]. Future studies could leverage the dual-process theories from psychology to explore the reverse phenomenon. As such, it becomes important for future studies of the mismatch to investigate whether the study method could impact how frequently and if the mismatch is observed.

Lastly, we seem to have vast evidence from our data that *people might not feel the need to express concerns since they already engage in privacy protective behaviors*. In fact, we saw evidence of related explanations, such as expressing being unconcerned because they do not expose themselves to the potentially concerning situation, or the situation not being significant, frequent, or disruptive enough. With this in mind, an umbrella explanation would be that people did not express being concerned because their risk valuation was not high enough for a concern to be expressed. Indeed, in many of the cases where we observed a mismatch between specific concern and behavior, there was no mismatch when we looked at the participant’s expressed preferences. This is of significance for the privacy research community to consider, as a significant portion of work in this area focuses on concern as its main proxy for privacy perspectives, and these studies often play a role in public policy discussions.

As noted in the Introduction, it may be instructive to tie the current results to those originating from the stream of research on the so-called privacy paradox. Experiments (and meta-analyses) of the traditional paradox have provided evidence of both *mismatches* between mental states (such as attitudes, concerns, or intentions) and privacy behaviors [7, 11, 25, 29], and of scenarios where mental states are *consistent* with privacy behaviors (see [17, 39] and, in the context of security behavior, [18]). The privacy-paradox literature highlights the importance of context in interpreting both self-reported privacy claims and actual behaviors [6]: a mismatch between privacy attitudes and behaviors does *not* (necessarily) imply that consumers do not care for privacy; similarly, a correspondence of attitudes and behaviors does not imply that consumers are *always* able to act in the marketplace in manners consistent with

their desired degree of privacy. This healthy degree of caution over simplistic applications of revealed-preferences economic argument in the realm of privacy applies also to our current analysis of a reverse privacy paradox. Ultimately, we believe that the results suggest caution in interpreting a self-reported lack of concern for privacy as evidence of consumers simply not caring for privacy. Rather, we interpret the results as reinforcing the need for nuanced analysis of individuals’ privacy stances (and behaviors) that takes into account both individual heterogeneity in mental models and significant heterogeneity across the scenarios where diverse privacy concerns may arise. In fact, as observed in Section 6.2, both general and specific mismatches did not imply that participants did not care for privacy, but rather that their actions and their concerns varied across scenarios. This lesson may be helpful to policy makers as well when they consider policy initiatives in the realm of privacy and attempt to assess the extent to which consumers *truly* care for privacy. The fact that the overwhelming majority of participants in our sample did engage in diverse privacy behaviors, including those participants who expressed indifferent or negative views towards privacy, should suggest skepticism of arguments that, in our digital age, privacy has become a social norm of the past [30].

Moving forward, it may be beneficial to be cautious about the use of language related to privacy concern for the reasons discussed here, as well as because survey questions designed to measure privacy concern may in fact be measuring other constructs [15]. In cases where it is of interest to researchers to specifically focus on concern, understanding what drives participants’ responses is key to correctly interpreting and using those results. Furthermore, while the Theory of Reasoned Action [19] has long included the importance of this feedback loop between behavior and perspectives, privacy literature seems to frequently focus on the connection between perspectives and behavior engagement. Moving forward, it would also be pertinent to investigate the reverse influence.

9 CONCLUSIONS AND FUTURE WORK

The inverse form of privacy mismatch this manuscript set out to explore did seem to appear within our sample of respondents. In a substantial portion of that sample, even participants who agreed with a characterization of themselves as someone who did not care or was not concerned about privacy frequently engaged in privacy-seeking behaviors. In fact, quite often, general mismatches did not arise because participants did not care *generally* for privacy; instead, their actions and their concerns varied across scenarios. Furthermore, both general and specific mismatches arose across participants for a multiplicity of reasons. Ultimately, many of our participants’ comments appear consistent with the notion within privacy scholarship that preferences *for* privacy should be understood within their context [28]—and so as well, it turns out, should preferences *against* privacy (such as claims of not caring or not being worried about it) be understood. Future work may focus on extending the analysis of the emergence of this type of attitudes/behaviors gap by considering more diverse and larger populations, and—building on the work presented here—complementing the survey-based analysis with experimental work.

ACKNOWLEDGMENTS

We thank the anonymous reviewers and shepherd for their suggestions for improving this manuscript. This research was supported in part by a Carnegie Mellon University CyLab partnership with NortonLifelock.

REFERENCES

- [1] Alessandro Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM conference on Electronic commerce*, pages 21–29, 2004.
- [2] Alessandro Acquisti, Idris Adjerid, Rebecca Balebako, Laura Brandimarte, Lorrie Faith Cranor, Saranga Komanduri, Pedro Giovanni Leon, Norman Sadeh, Florian Schaub, Manya Sleeper, et al. Nudges for privacy and security: Understanding and assisting users' choices online. *ACM Computing Surveys (CSUR)*, 50(3):1–41, 2017.
- [3] Alessandro Acquisti, Idris Adjerid, and Laura Brandimarte. Gone in 15 seconds: The limits of privacy transparency and control. *IEEE Security and Privacy*, 11(4):72–74, July 2013.
- [4] Alessandro Acquisti, Laura Brandimarte, and Jeff Hancock. How privacy's past may shape its future. *Science*, 375(6578):270–272, 2022.
- [5] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.
- [6] Alessandro Acquisti, Laura Brandimarte, and George Loewenstein. Secrets and Likes: The Drive for Privacy and the Difficulty of Achieving It in the Digital Age. *Journal of Consumer Psychology*, 30(4):736–758, oct 2020.
- [7] Alessandro Acquisti and Ralph Gross. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *International workshop on privacy enhancing technologies*, pages 36–58. Springer, 2006.
- [8] Alessandro Acquisti and Jens Grossklags. Privacy and Rationality in Individual Decision Making. *IEEE Security and Privacy*, 3(1):26–33, jan 2005.
- [9] Michael Adorjan and Rosemary Ricciardelli. A new privacy paradox? Youth agentic practices of privacy management despite “nothing to hide” online. *Canadian Review of Sociology*, 56(1):8–29, 2019.
- [10] Irwin Altman. *The environment and social behavior: privacy, personal space, territory, and crowding*. ERIC, 1975.
- [11] Susanne Barth, Menno DT de Jong, Marianne Junger, Pieter H Hartel, and Janina C Roppelt. Putting the privacy paradox to the test: Online privacy and security behaviors among users with technical knowledge, privacy awareness, and financial resources. *Telematics and informatics*, 41:55–69, 2019.
- [12] F. Z. Borgesius. Informed consent: We can do better to defend privacy. *IEEE Security and Privacy*, 13(2):103–107, 2015.
- [13] Long Chen, Yadong Huang, Shumiao Ouyang, and Wei Xiong. The data privacy paradox and digital demand. Technical report, National Bureau of Economic Research, 2021.
- [14] Noam Cohen. Law students teach scalia about privacy and the web. <https://www.nytimes.com/2009/05/18/technology/internet/18link.html>, 2009. Accessed Feb, 2022.
- [15] Jessica Colnago, Lorrie Faith Cranor, Alessandro Acquisti, and Kate Hazel Jain. Is it a concern or a preference? An investigation into the ability of privacy scales to capture and distinguish granular privacy constructs. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 331–346, Boston, MA, August 2022. USENIX Association.
- [16] Jessica Colnago, Yuanyuan Feng, Tharangini Palanivel, Sarah Pearman, Megan Ung, Alessandro Acquisti, Lorrie Faith Cranor, and Norman Sadeh. Informing the Design of a Personalized Privacy Assistant for the Internet of Things. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems*, pages 1–13, 2020.
- [17] Tobias Dienlin, Philipp K Masur, and Sabine Trepte. A longitudinal analysis of the privacy paradox. *New Media & Society*, page 14614448211016316, 2019.
- [18] Serge Egelman, Marian Harbach, and Eyal Peer. Behavior ever follows intention? a validation of the security behavior intentions scale (SeBIS). In *Proceedings of the 2016 CHI conference on human factors in computing systems*, pages 5257–5261, 2016.
- [19] Martin Fishbein and Icek Ajzen. Introduction. In *Belief, Attitude, Intention, and Behavior: An Introduction to Theory and Research*, chapter 1. Addison-Wesley, 1975.
- [20] Reza Ghaiumy Anaraky, Kaileigh Angela Byrne, Pamela J. Wisniewski, Xinru Page, and Bart Knijnenburg. To disclose or not to disclose: Examining the privacy decision-making processes of older vs. younger adults. In *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, pages 1–14, 2021.
- [21] Hongxia Jin, G. Saldamli, R. Chow, and B. P. Knijnenburg. Recommendations-based location privacy control. In *2013 IEEE International Conference on Pervasive Computing and Communications Workshops (PERCOM Workshops)*, pages 401–404. IEEE, March 2013.
- [22] Ruogu Kang, Stephanie Brown, Laura Dabbish, and Sara Kiesler. Privacy attitudes of Mechanical Turk workers and the U.S. public. In *10th Symposium On Usable Privacy and Security (SOUPS 2014)*, pages 37–49, Menlo Park, CA, jul 2014. USENIX Association.
- [23] Spyros Kokolakis. Privacy attitudes and privacy behaviour: A review of current research on the privacy paradox phenomenon. *Computers and Security*, 64:122–134, jan 2017.
- [24] Scott Lederer, Jason I. Hong, Anind K. Dey, and James A. Landay. Personal privacy through understanding and action: Five pitfalls for designers. *Personal and Ubiquitous Computing*, 8(6):440–454, 2004.
- [25] Michelle Madejski, Maritza Johnson, and Steven M Bellovin. A study of privacy settings errors in an online social network. In *2012 IEEE international conference on pervasive computing and communications workshops*, pages 340–345. IEEE, 2012.
- [26] Caroline Lancelot Miltgen and Dominique Peyrat-Guillard. Cultural and generational influences on privacy concerns: A qualitative study in seven European countries. *European journal of information systems*, 23(2):103–125, 2014.
- [27] Mercury News. Mark Zuckerberg's building a wall, and his neighbors don't like it. <https://www.mercurynews.com/2016/06/29/mark-zuckerbergs-building-a-wall-and-his-neighbors-dont-like-it>, 2016. Accessed Feb, 2022.
- [28] Helen Nissenbaum. Privacy as contextual integrity. *Wash. L. Rev.*, pages 101–139, 2004.
- [29] Patricia A. Norberg, Daniel R. Horne, and David A. Horne. The privacy paradox: Personal information disclosure intentions versus behaviors. *Journal of Consumer Affairs*, 41(1):100–126, 2007.
- [30] Ian Paul. Facebook CEO challenges the social norm of privacy. PCWorld, 2009. Accessed Feb, 2022.
- [31] Darhl M. Pedersen. Model for types of privacy by privacy functions. *Journal of environmental psychology*, 19(4):397–405, 1999.
- [32] Sandra Petronio. Communication boundary management: A theoretical model of managing disclosure of private information between marital couples. *Communication Theory*, 1(4):311–335, November 1991.
- [33] Chanda Phelan, Cliff Lampe, and Paul Resnick. It's creepy, but it doesn't bother me. In *Proceedings of the 2016 CHI Conference on Human Factors in Computing Systems*, CHI '16, pages 5240–5251. ACM, 2016.
- [34] Sören Preibusch. Guide to measuring privacy concern: Review of survey and observational instruments. *International Journal of Human-Computer Studies*, 71(12):1133–1143, 2013.
- [35] Daniel J. Solove. A Taxonomy of Privacy. *University of Pennsylvania Law Review*, 154(3), 2006.
- [36] Daniel J Solove. The myth of the privacy paradox. *Geo. Wash. L. Rev.*, 89, 2021.
- [37] Qizhang Sun, Martijn C. Willemsen, and Bart P. Knijnenburg. Unpacking the intention-behavior gap in privacy decision making for the internet of things (IoT) using aspect listing. *Computers & Security*, 97:101924, 2020.
- [38] Jenny Tang, Eleanor Birrell, and Ada Lerner. Replication: How well do my results generalize now? The external validity of online privacy and security surveys. In *Eighteenth Symposium on Usable Privacy and Security (SOUPS 2022)*, pages 367–385, Boston, MA, August 2022. USENIX Association.
- [39] Sabine Trepte, Michael Scharkow, and Tobias Dienlin. The privacy calculus contextualized: The influence of affordances. *Computers in Human Behavior*, 104:106115, 2020.
- [40] Samuel Warren and Louis Brandeis. The Right to Privacy. *Harvard Law Review*, 4(5):193–220, 1890.
- [41] Benjamin Wittes and Emma Kohse. *The privacy paradox II: Measuring the privacy benefits of privacy threats*. Center for Technology Innovation at Brookings, 2017.

#	Behavior	Specific concerns	Specific preferences
1	Used an alias (fake name) or dummy account when I did not want my real identity tied to an activity, for example, sharing my art, photos, or writing, commenting on public forums, etc	I worry about my real identity being tied to some of my activities	I do not want my real identity to be tied to some of my activities
2	Changed the status indicator on an app or website to not show I was online	I worry about others being able to see when I am online	I want to have control over when others can see when I am online
3	Chose to go to certain places by myself to keep my activities private	I worry about others knowing about my activities in certain places	I do not want others to find out about my activities in certain places
4	Made sure I was alone when using the internet for private matters like gift shopping, private emails, porn, etc.	I worry about others seeing how I use the internet for private matters	I want my use of the internet for private matters to remain unknown to others
5	Refused to answer when others tried to contact me so I could have some private time	I worry about having my private time interrupted	I want to be able to have private time
6	Used a different communication service (app, tool, etc) than I normally do for conversations that I wanted to keep private	I worry about others finding out about conversations that I want to keep private	I want to keep some conversations private
7	Used a secondary email for specific services like purchases, newsletter, reviews, subscriptions to particular sites, etc	I worry about sharing my email address	I want to reduce how much I share my email address
8	Used specific credit cards for purchases that I wanted to keep private	I worry about not being able to keep some purchases private	I want to keep some purchases private
9	Changed the subject or set boundaries when the conversation turned to sensitive topics I preferred not to discuss	I worry about discussing sensitive topics	I do not want to discuss certain topics
10	Selectively chose which phone calls to pick up.	I worry about unwanted phone calls	I do not want to receive unwanted phone calls
11	Stopped using apps or websites because of privacy concerns about them or the company that owns them	I worry about whether companies that own the apps and websites I use protect my privacy	I want to have my privacy protected when I use apps and websites
12	Complained to companies when I thought they had collected or sold my information without my permission	I worry about my information being collected or sold without my permission	I want to have control over my information being collected or sold
13	Reprimanded someone for behavior that invaded my privacy, such as entering my private room or looking through my stuff	I worry about others invading my private spaces	I want to keep my private spaces private
14	Tried to prevent others from seeing what's inside my wallet or concealed my actions when conducting transactions at an ATM	I worry about others learning about my financial details	I do not want others learning about my financial details
15	Considered whether a company, app, or service may be harvesting my personal information before using their services	I worry about my information being harvested	I do not want to have my personal information harvested
16	Considered whether the benefit of signing up for a newsletter or using an app was worth the collection of my personal data	I worry about the trade-off of providing my personal data in exchange for a service	I want to have a good trade-off between providing my personal data and the service I am using
17	Took steps to not be tracked online, for example, by deleting cookies, leaving a website that requested cookies, or by changing settings that would allow websites to track me	I worry about being tracked online	I do not want to be tracked online
18	Checked the options for storing cookies and data collection for a website and opted-out of the ones I did not support	I worry about not being able to opt-out of data collection practices that I do not agree with	I want to be able to opt-out of data collection practices that I do not agree with
19	Avoided personality quizzes and other features that could be used to harvest information from social profiles	I worry about information from my social media being collected through personality quizzes and other tools	I do not want my information from my social media to be collected through personality quizzes and other tools
20	Attempted to hide my internet footprint, for example, by deleting information about me	I worry about all of the information about me available online	I want to have less information about me available online
21	Locked my devices, using a password, fingerprint, etc, or kept my devices with me to avoid others gaining access to them	I worry about others gaining access to my devices	I want to limit who has access to my devices
22	Took steps so that others couldn't see what I was doing, for example, minimized my screen, sat with my back to a wall, shielded my phone or documents, used a screen protector, etc	I worry about others seeing what I am doing	I want to keep some of my activities private
23	Declined to provide sensitive information, like banking information or my SSN, when it wasn't clear that the request was legitimate or necessary	I worry about providing sensitive information when I'm not sure the request is legitimate or necessary	I do not want to provide sensitive information when I am not sure the request is legitimate or necessary
24	Gave the least amount of information when answering a personal question	I worry about sharing too much information when answering personal questions	I do not want to share too much information when answering personal questions
25	Refused to answer surveys that required personal information	I worry about sharing too much information when answering surveys	I do not want to share too much information when answering surveys
26	Was careful about what I said and what I let other people know about me	I worry about others learning too much about me	I want to limit how much others can learn about me
27	Requested to be put on the Do Not Call list	I worry about unwanted phone calls	I do not want to receive unwanted phone calls
28	Secured or shredded sensitive documents	I worry about others gaining access to my sensitive documents	I want to keep sensitive documents safe from others

Table 6 continued from previous page

#	Behavior	Specific concerns	Specific preferences
29	Made sure sites that hold sensitive information were encrypted (had the lock icon)	I worry about sites with sensitive information about me not being secure	I want sites that hold sensitive information to be encrypted and secure
30	Only disclosed financial information with companies that are well established and reputable	I worry about sharing financial information	I want to only share financial information with companies I trust
31	Avoided using government websites to prevent government surveillance	I worry about government surveillance	I do not want to be under government surveillance
32	Reduced or censored my activity online so that the government or other entities could not spy on me	I worry about being spied on online	I do not want to be spied on online
33	Took steps to minimize the chance of being hacked like using an adblocker, avoiding the dark web, or using a firewall or anti-virus	I worry about being hacked	I do not want to be hacked
34	Cleared my browsing history	I worry about others finding out about what websites I visit	I want to keep the websites I visit private
35	Went inside when I noticed someone peering into my backyard where I was gathered with friends and/or family	I worry about strangers looking in my backyard	I do not want to be observed by strangers when I'm in my backyard
36	Provided intentionally incorrect information to strangers when asked for personal information	I worry about providing sensitive information when I'm not sure the request is legitimate or necessary	I do not want to provide correct personal information to strangers
37	Provided intentionally incorrect personal details about me when creating an account online	I worry about providing sensitive information when I'm not sure the request is legitimate or necessary	I do not want to provide correct personal information when creating some accounts online
38	Took steps so that services, apps, or websites couldn't access my location without my knowledge, for example, turned my location off, or kept it off, turned my phone off, etc	I worry about others gaining access to my location data	I want to keep my location information private
39	Asked for information to be emailed to avoid providing my address or phone number	I worry about others gaining access to my address or phone number	I want to keep my address or phone number private
40	Avoided shopping online or chose to shop in-person due to privacy concerns	I worry about my privacy when shopping online	I prefer to shop in-person rather than online
41	Changed from an online to an offline method for sensitive activities like banking, sensitive documents, or private conversations	I worry about performing sensitive activities online	I prefer to perform sensitive activities offline rather than online
42	Chose to pay using a payment method that would track me less than a credit card, for example, using cash or PayPal	I worry about being tracked through my credit card	I prefer to pay using a payment method that won't track me rather than my credit card
43	Opted for a different method of authentication other than facial recognition	I worry about facial recognition	I do not want to use facial recognition technology
44	Used login and password best practices, like not sharing my password, logging out of accounts, using strong passwords, etc	I worry about the security of my accounts	I want to keep my accounts secure
45	Checked the permissions needed by an app when downloading it or the permissions already granted to installed apps, for example, which apps can access my camera or location	I worry about what access mobile apps have to my device and information	I want to know what data and information mobile apps need and already have from my device
46	Denied permissions on mobile apps or websites, like sharing location or access to the camera, when I didn't agree with them	I worry about what access mobile apps have to my device and information	I want to have control over the permissions required by apps and websites
47	Covered my devices' cameras	I worry about others seeing me through my camera without my knowledge	I do not want others to see me through my camera without my knowledge
48	Protected the privacy of my home by closing curtains or doors and installing fences, etc	I worry about others observing me while I'm at home	I want to keep my home private
49	Covered myself so that others would not see my body, for example, with a towel at the pool or when coming out of the shower, the dressing room curtain at a store, etc.	I worry about others seeing my exposed body	I do not want others seeing my exposed body
50	Physically removed myself from situations where I didn't want to be seen or answer questions.	I worry about being in situations where I would not want to be seen or answer questions	There are times I do not want to be seen or answer questions
51	Read the privacy policy for a service or tool I was going to use	I worry about how services and tools that I use protect my privacy	I want to know about the privacy practices of the services and tools I use
52	Ignored links and pop-ups from websites that try to collect personal information like sign-up or newsletter pop-ups	I worry about websites collecting personal information	I do not want to share personal information with websites
53	Put up a front so people couldn't know the real me or how I felt	I worry about others finding out too much about me or how I feel	I do not want others finding out too much about me or how I feel
54	Set up protection services, for example, identity theft protection, security systems, etc	I worry about my security	I want to have security
55	Wore sunglasses, a hat, or other strategies to avoid being recognized	I worry about being recognized	There are times I do not want to be recognized

Table 6 continued from previous page

#	Behavior	Specific concerns	Specific preferences
56	Did not use an app or service, or deleted an app, because it asked for too much information or did not allow me to opt-out of data collection	I worry about apps that collect too much information or do not allow me to opt-out of the data collection	I do not want apps and services to collect too much of my information
57	Reported suspicious activity, for example, unknown charges to account, identity theft, potential scams or spam, etc	I worry about suspicious activities like scams	I do not want to be the victim of suspicious activities
58	Asked friends not to tag me on their posts	I worry about my friends tagging me on their posts	I do not want to be tagged on my friends' posts
59	Sought or told others that I needed to be left alone	I worry about not being able to be left alone	I sometimes want to be alone
60	Avoided letting stores save my information like credit card details, address, or password	I worry about stores saving my information	I do not want stores to save my information
61	Avoided searching for certain topics online for fear of being monitored	I worry about my online searches being monitored	I do not want to have my online searches monitored
62	Refused to use social media logins on websites	I worry about using social media logins on websites	I do not want to use social media logins on websites
63	Changed settings or blocked people on social media to only allow certain people to view my information (posts, photos, etc.)	I worry about who can see my information on social media	I want to limit who can see my information on social media
64	Chose not to have, deleted, or reduced my use of social media accounts due to privacy concerns	I worry about my privacy in social media	I do not want information about me to be revealed through my use of social media
65	Deleted old posts, photos, or information from social media platforms to protect my privacy	I worry about my privacy in social media	I want to have control over what information is available about me on my social media accounts
66	Spoke more quietly, moved away, or wrote something down instead of saying it to avoid others from overhearing a sensitive conversation	I worry about others overhearing me when having a sensitive conversation	I want to be able to have sensitive conversations without being overheard
67	Tried to avoid security cameras in public areas	I worry about security cameras in public areas	I do not want to be recorded in security camera footage while in public I do not want my phone to spy on me
68	Kept my phone away from me when having private conversations to prevent it from spying on what was being said	I worry about my phone spying on me	
69	Chose to use services that are more privacy protective like the Tor browser, DuckDuckGo, Signal app, etc	I worry about my privacy online	I want to have privacy online
70	Used tools to increase my privacy like private browsing or incognito mode, blocking third-party cookies on my browser, encrypted emails, VPN, adblockers, etc	I worry about my privacy online	I want to have privacy online
71	Only used commonly used websites to make sure my data is protected	I worry about using unknown websites	I want to keep my data protected online
72	Avoided using public or unsecured WiFi	I worry about using public or unsecured WiFi	I want my browsing and data to be safe when using WiFi
73	Added a password to my home WiFi or phone hotspot	I worry about others accessing my home WiFi or phone hotspot	I want to keep my home WiFi or phone hotspot protected
74	Chose to shop online instead of in-person due to privacy concerns	I worry about my privacy when shopping in person	I prefer to shop online rather than in-person
75	Encrypted or password protected sensitive files	I worry about others gaining access to sensitive files	I want to keep sensitive files secured
76	Used different emails for work and personal affairs	I worry about having no separation between my work and personal life	I want to have separation between my work and personal life
77	Set email filters and spam filters to avoid receiving unwanted emails	I worry about receiving unwanted emails	I do not want to receive unwanted emails

Table 6: Behaviors and associated specific preferences and concerns (Behavior Elicitation Study, Study 1)

Behavior	Motivation								Frequency						
	M1	M2	M3	M4	M5	M6	M7	M8	F1	F2	F3	F4	F5	F6	F7
B1	6	0	16	1	1	0	0	9	8	5	8	6	4	1	1
B2	2	2	19	2	1	0	2	9	9	5	11	5	5	2	0
B3	0	1	14	2	3	3	0	14	14	9	10	0	2	2	0
B4	0	1	27	0	2	3	0	5	4	1	10	7	7	8	1
B5	1	1	9	12	7	0	5	1	1	9	15	9	2	0	0
B6	4	0	13	1	1	2	0	16	15	5	11	4	1	0	1
B7	2	0	8	12	5	4	2	4	4	0	6	9	5	13	0
B8	1	1	6	2	2	3	0	17	15	8	4	2	1	0	2
B9	5	7	9	4	6	0	0	4	4	3	20	4	4	0	0
B10	3	1	8	13	4	3	3	0	0	3	7	7	4	14	0
B11	3	1	11	0	2	10	0	10	9	4	18	1	1	3	1
B12	0	0	5	0	1	2	0	28	28	7	1	0	0	0	0
B13	1	2	22	0	4	1	0	4	4	13	12	1	2	2	0
B14	11	0	6	0	2	14	0	4	4	5	9	8	6	5	0
B15	5	1	16	0	2	6	0	5	4	3	9	10	6	2	1
B16	1	0	16	4	1	7	0	6	6	2	12	7	4	4	0
B17	2	0	16	0	1	12	0	6	5	9	12	6	3	1	1
B18	1	2	18	0	4	5	0	7	7	10	11	2	5	2	0
B19	5	0	14	1	0	5	2	10	8	4	5	4	7	7	2
B20	6	0	18	0	1	3	1	6	4	9	17	1	2	0	2
B21	4	0	12	1	2	16	0	0	0	1	3	2	9	20	0
B22	1	1	23	1	1	7	0	3	3	5	17	8	2	2	0
B23	8	0	2	1	0	19	0	3	2	0	9	4	3	14	1
B24	5	1	24	1	1	3	0	0	0	6	18	5	5	1	0
B25	6	1	16	0	1	6	1	4	3	9	16	2	3	1	1
B26	7	8	17	0	4	0	0	0	0	5	6	12	7	6	0
B27	1	2	7	4	7	2	1	12	10	2	5	4	5	8	2
B28	7	1	6	0	1	16	0	5	5	5	8	7	6	5	0
B29	6	0	5	1	0	15	0	10	8	5	10	4	5	3	2
B30	9	0	5	0	3	18	0	2	0	1	4	4	11	15	2
B31	1	0	4	1	1	3	0	25	24	7	2	1	0	0	1
B32	2	0	12	1	0	2	0	18	17	6	5	4	1	1	1
B33	10	0	4	0	1	16	0	4	3	1	6	6	8	10	1
B34	0	1	16	4	8	2	2	3	3	8	16	8	1	0	0
B35	4	0	4	0	1	0	0	25	21	5	1	2	0	1	4
B36	17	1	7	2	0	5	0	5	4	10	14	3	4	1	1
B37	5	1	13	4	2	5	0	5	3	9	16	5	0	0	2
B38	7	0	21	1	0	6	0	1	1	7	9	7	6	6	0
B39	7	0	6	3	1	3	1	13	11	4	9	3	2	3	2
B40	2	0	4	6	2	6	0	16	16	14	5	1	0	0	0
B41	1	1	6	1	2	8	0	17	16	7	8	2	2	0	1
B42	2	0	6	7	4	4	0	14	14	5	9	7	2	0	0
B43	1	0	10	3	2	4	2	13	8	1	5	1	5	10	5
B44	12	0	5	0	1	18	1	0	0	0	6	5	12	14	0
B45	4	0	16	0	1	9	0	6	6	2	11	2	8	7	0
B46	6	1	16	1	1	12	0	1	1	3	10	9	3	12	0
B47	8	0	12	0	0	5	0	10	9	3	10	1	6	5	1
B48	13	1	15	0	2	3	0	1	1	2	9	8	8	7	0
B49	2	12	14	0	2	0	2	4	3	2	7	5	8	10	1
B50	7	1	10	9	7	0	0	3	2	5	17	5	4	3	1
B51	5	1	6	4	6	5	1	7	7	15	10	2	1	0	0
B52	7	1	8	5	3	8	0	2	2	0	1	5	11	15	0
B53	2	11	9	2	7	1	0	4	4	8	13	7	1	3	0
B54	8	0	2	0	0	12	0	14	14	5	10	1	2	4	0
B55	0	2	5	2	2	0	1	25	25	10	2	0	0	0	0
B56	5	0	16	0	0	11	0	5	3	8	10	6	4	4	2
B57	5	2	1	1	1	16	0	7	6	11	8	1	0	6	1
B58	1	0	16	0	1	0	1	17	16	7	6	1	2	3	1
B59	1	2	12	6	7	0	4	3	3	10	13	7	2	0	0
B60	5	0	4	1	2	18	0	3	3	4	11	4	8	3	0
B61	3	2	7	0	1	3	0	18	17	10	5	0	1	0	1
B62	4	1	14	2	3	6	0	7	6	7	9	7	1	6	1
B63	4	1	22	3	3	0	1	2	2	4	17	6	3	4	0
B64	5	1	9	4	7	0	0	9	8	10	12	0	2	2	1
B65	5	4	16	0	5	1	1	6	5	8	20	1	1	2	1
B66	0	3	19	0	3	4	1	7	6	4	15	9	2	0	1
B67	1	0	5	1	1	2	0	26	26	7	3	0	0	0	0
B68	0	1	7	1	1	1	0	27	27	8	2	1	0	0	0
B69	2	1	10	0	1	2	0	18	17	7	5	2	1	1	1
B70	6	0	19	0	1	7	0	1	1	7	9	8	7	2	0
B71	8	0	6	2	2	11	1	5	4	4	1	11	8	6	1
B72	3	1	5	4	2	12	0	8	7	9	4	7	5	2	1
B73	6	1	4	1	3	19	0	1	1	4	2	1	2	25	0
B74	2	4	4	6	0	1	0	20	20	5	5	3	3	1	0
B75	1	0	11	0	0	14	0	10	8	11	7	2	3	3	2

Table 7 – Continued from previous page

Behavior	Motivation								Frequency						
	M1	M2	M3	M4	M5	M6	M7	M8	F1	F2	F3	F4	F5	F6	F7
B76	1	3	9	4	13	4	0	2	1	1	3	7	8	15	1
B77	1	0	1	15	8	6	0	5	5	3	12	6	5	5	0

Table 7: Distribution of frequency and motivation selections across all behaviors (Study 1A). Motivation legend: M1-Safety, M2-Social norms, M3-Privacy, M4-Convenience, M5-Practical reasons, M6-Security, M7-Another reason, M8-I don't do this. Caption legend: F1-Never, F2-Rarely, F3-Sometimes, F4-Frequently, F5-Most of the time, F6-Every time the opportunity presents itself, F7-Not applicable.

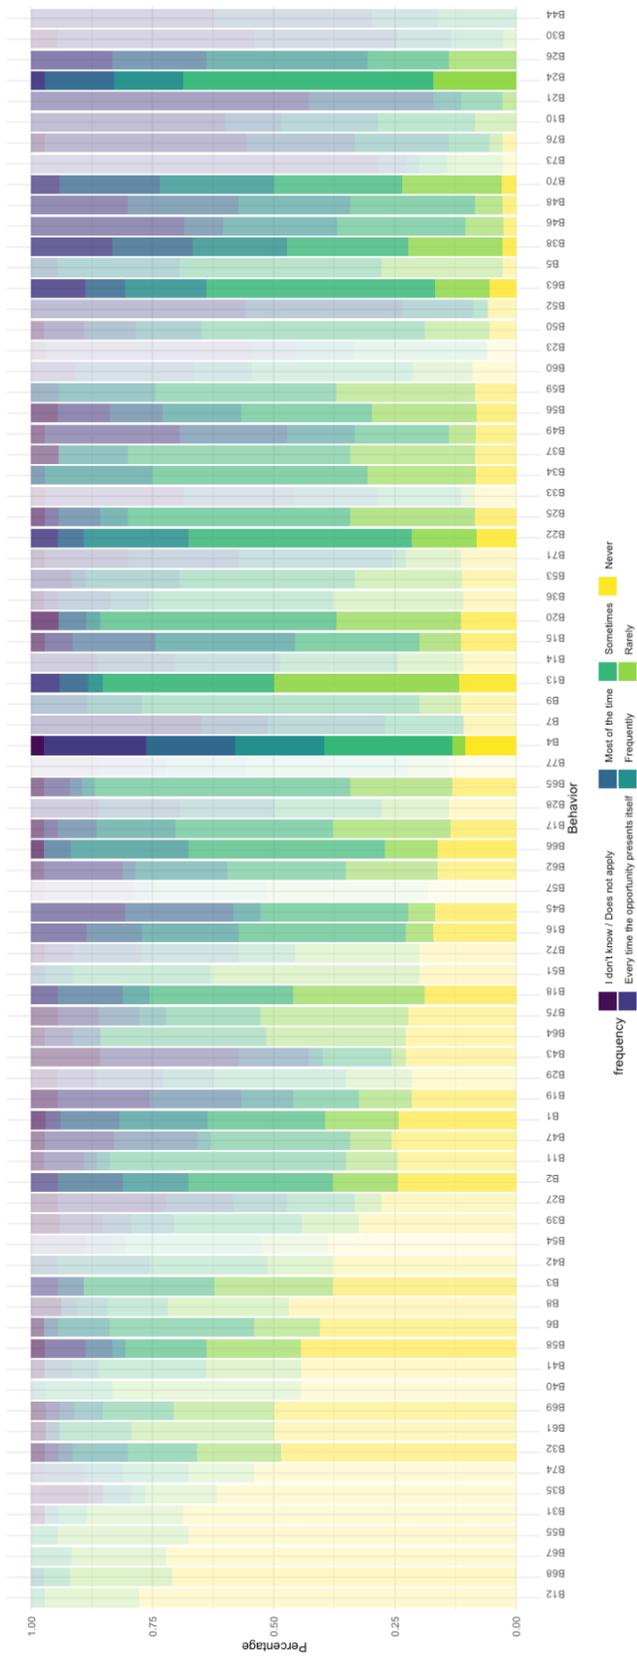


Figure 1: Distribution of behavior engagement frequency per participant, where lighter-colored bars represent behaviors that were less frequently engaged in due to a privacy motivation (Study 1A).

Behavior	Specific concern					Specific preference				
	Strongly disagree (1)	Somewhat disagree (2)	Neither agree nor disagree (3)	Somewhat agree (4)	Strongly agree (5)	Strongly disagree (1)	Somewhat disagree (2)	Neither agree nor disagree (3)	Somewhat agree (4)	Strongly agree (5)
B1	1	0	1	5	2	0	1	0	4	4
B2	0	3	1	3	3	0	0	2	1	7
B3	1	2	1	2	2	1	3	2	2	0
B4	1	0	1	2	4	0	1	2	1	4
B5	1	2	1	5	2	0	0	0	2	9
B6	0	1	2	3	2	0	0	0	1	7
B7	0	2	3	9	2	0	1	2	6	7
B8	0	0	2	2	1	0	0	1	1	3
B9	2	1	2	6	0	0	0	4	5	2
B10	1	2	1	1	3	0	0	0	3	5
B11	0	1	0	2	4	0	0	0	1	6
B12	0	0	0	5	5	0	0	1	1	8
B13	0	0	1	3	4	0	0	0	0	8
B14	0	0	0	4	3	0	0	0	2	5
B15	0	1	1	3	2	0	0	0	2	5
B16	2	0	1	2	3	0	0	2	2	4
B17	0	1	1	4	2	0	0	0	2	6
B18	0	0	2	6	2	0	0	1	3	6
B19	0	1	2	5	6	0	0	2	4	8
B20	0	0	4	3	3	0	0	2	5	3
B21	0	0	1	9	1	0	0	1	4	6
B22	0	0	5	4	0	0	0	0	2	7
B23	0	0	1	2	5	0	0	1	0	7
B24	0	0	3	7	0	0	0	1	5	4
B25	0	2	1	6	1	0	2	1	6	1
B26	0	0	3	5	2	0	0	2	5	3
B27	1	0	1	6	4	0	0	0	2	10
B28	0	1	0	3	4	0	0	1	1	6
B29	0	0	0	6	3	0	0	0	2	7
B30	0	0	0	3	6	0	0	1	2	6
B31	2	3	1	4	0	0	0	0	6	4
B32	0	2	1	7	2	0	0	2	3	7
B33	0	0	0	4	7	0	0	0	1	10
B34	0	1	3	3	1	0	0	0	4	4
B35	2	0	2	2	2	0	0	1	1	6
B36	0	0	1	3	6	0	2	2	4	2
B37	0	0	0	5	10	2	1	3	6	3
B38	0	1	4	2	4	0	0	2	5	4
B39	0	1	1	3	1	0	1	1	0	4
B40	0	3	2	4	1	0	3	4	1	2
B41	0	0	4	3	3	0	0	3	4	3
B42	0	5	4	3	0	3	1	4	1	3
B43	0	0	2	4	1	0	2	0	2	3
B44	0	0	0	5	3	0	0	0	2	6
B45	0	0	1	5	1	0	0	1	3	3
B46	0	1	1	3	5	0	0	0	3	7
B47	0	2	1	1	2	0	0	0	0	6
B48	1	3	1	2	4	0	0	1	0	10
B49	1	1	3	3	1	0	0	3	2	4
B50	0	0	6	7	1	0	1	0	7	6
B51	0	0	1	5	2	0	0	1	5	2
B52	0	0	4	4	1	0	0	1	4	4
B53	0	1	3	3	1	0	1	2	4	1
B54	0	1	4	3	3	0	0	0	3	8
B55	1	4	4	2	0	0	0	3	6	2
B56	0	0	2	2	3	0	0	1	2	4
B57	0	0	2	4	3	0	0	1	1	7
B58	3	3	4	2	0	1	5	3	1	2
B59	2	1	0	1	1	0	0	0	1	4
B60	0	0	3	2	4	0	0	0	4	5
B61	0	1	0	2	1	0	0	0	2	2
B62	1	1	2	4	4	0	0	5	3	4
B63	0	1	3	4	5	0	0	1	1	11
B64	0	1	2	4	6	0	0	3	7	3
B65	0	1	1	8	1	0	0	1	5	5
B66	0	0	1	5	2	0	0	0	1	7
B67	3	3	1	1	0	1	4	1	2	0
B68	1	3	0	8	2	0	0	0	3	11
B69	0	1	1	8	4	0	0	1	4	9
B70	0	0	1	5	3	0	0	0	3	6
B71	0	0	0	1	6	0	0	0	1	6
B72	0	3	2	2	2	0	0	1	1	7

Table 8 continued from previous page

Behavior	Specific concern					Specific preference				
	Strongly disagree (1)	Somewhat (disagree) (2)	Neither agree nor disagree (3)	Somewhat agree (4)	Strongly agree (5)	Strongly disagree (1)	Somewhat (disagree) (2)	Neither agree nor disagree (3)	Somewhat agree (4)	Strongly agree (5)
B73	1	1	1	2	1	0	0	1	1	4
B74	4	5	2	0	0	0	0	6	3	2
B75	0	0	0	1	7	0	0	0	0	8
B76	1	2	2	2	3	0	0	1	2	7
B77	0	4	2	2	2	0	0	0	3	7

Table 8: Distribution of agreement across all behaviors for the matching specific statements of concern and preference (Study 1B).