Would a privacy fundamentalist sell their DNA for $1000... if nothing bad happened as a result? The Westin categories, behavioral intentions, and consequences

Allison Woodruff
Google
1600 Amphitheatre Pkwy
Mountain View, CA 94043
woodruff@acm.org

Vasyl Pihur
Google
1600 Amphitheatre Pkwy
Mountain View, CA 94043
vpihur@google.com

Sunny Consolvo
Google
1600 Amphitheatre Pkwy
Mountain View, CA 94043
sconsolvo@google.com

Lauren Schmidt
Google
1600 Amphitheatre Pkwy
Mountain View, CA 94043
schmidtl@google.com

Laura Brandimarte
Carnegie Mellon University
5000 Forbes Av. HBH 2105C
Pittsburgh, PA 15213
lbrandim@andrew.cmu.edu

Alessandro Acquisti
Carnegie Mellon University
5000 Forbes Av. HBH 2105C
Pittsburgh, PA 15213
acquisti@andrew.cmu.edu

ABSTRACT

Westin's Privacy Segmentation Index has been widely used to measure privacy attitudes and categorize individuals into three privacy groups: fundamentalists, pragmatists, and unconcerned. Previous research has failed to establish a robust correlation between the Westin categories and actual or intended behaviors. Unexplored however is the connection between the Westin categories and individuals' responses to the *consequences* of privacy behaviors. We use a survey of 884 Amazon Mechanical Turk participants to investigate the relationship between the Westin Privacy Segmentation Index and attitudes and behavioral intentions for both privacy-sensitive scenarios and privacy-sensitive consequences. Our results indicate a lack of correlation between the Westin categories and behavioral intent, as well as a lack of correlation between the Westin categories and consequences. We discuss potential implications of this attitude-consequence gap.

1. INTRODUCTION

Privacy research pioneer Alan Westin conducted over thirty privacy-related surveys between 1978 and 2004 [25]. During this time, he developed a Privacy Segmentation Index consisting of three questions and a set of rules to translate participants' responses into three categories (fundamentalists, pragmatists, and unconcerned) [24, 25]. This index captures general privacy attitudes about consumer control, business, and laws and regulations. It has been hugely influential in the debate over privacy attitudes, and has been deployed by researchers in numerous studies, e.g., [13, 23, 26, 29].

Symposium on Usable Privacy and Security (SOUPS) 2014, July 9–11, 2014, Menlo Park, CA.

Nonetheless, concerns have long existed regarding the predictive power of Westin's categories and the assumptions underlying his Privacy Segmentation Index. First, previous research has failed to establish a significant correlation between the Westin categories (which capture broad, generic privacy attitudes) and context-specific, privacy-related behaviors, either actual or intended [13, 23, 29]. Second, researchers have raised concerns regarding unstated assumptions underlying the index, which presumes individuals make privacy decisions that are highly rational, reflective, and informed [42]. Instead, scholars have posited that incomplete information or decision-making biases, among other factors, may cause a gap between the general attitudes captured by the Westin categories and actual, specific privacy behavior [4]. Third, the instrument has not been updated since approximately 1995, and it is not obvious that it remains current in our Internet-centric world.

It is perhaps unsurprising that generic attitudes (such as those captured by Westin's Privacy Segmentation Index) are poor predictors of context-specific behaviors [15]. The so-called privacy paradox is often interpreted as the apparent lack of correlation between privacy attitudes and behaviors, and much work on this topic has focused on contrasting generic attitudes with hypothetical or observed behavior. However, one might suppose that general attitudes would be more successful at predicting responses to consequences. For example, one might imagine that a fundamentalist would object more strongly than an unconcerned to a personal photo being distributed widely on the Internet. In this manuscript we test the relationship between the Westin categories and a diverse, large set of scenarios, and examine the previously unexplored connection between those categories and individuals' reactions to privacy-relevant outcomes from those scenarios. In other words, we examine whether generic privacy attitudes are correlated with individuals' attitudes and behavioral intentions when hypothetical but specific consequences arising from the protection or disclosure of personal information are described. We survey 884 Amazon Mechanical Turk participants to investigate this relationship.

Supporting but extending previous literature, our results suggest a lack of correlation between the Westin categories and any of the scenarios we designed, independent of the type of data, actions, and context presented to the participants. Expanding previous literature, our results also suggest a lack of correlation between the Westin categories and actual outcomes, regardless of the material consequences associated with the disclosure of personal data. We discuss the potential implications of this apparent attitude-consequence gap for the motives and rationales underlying privacy attitudes, and for the design and evaluation of privacy "personas" or privacy segmentations.

Additionally, we explore several potential improvements to the Westin Privacy Segmentation Index. First, we report on a data-driven segmentation of responses to the Westin questions, which did not result in significantly better response prediction than the Westin categories. Second, we explore the implications of making the Westin questions more specific by replacing generic companies with particular brands; our results indicate this manipulation tends to make participants less privacy-sensitive. Third, we investigate whether other specific variables such as personality traits and demographics are more predictive of responses to scenarios or outcomes than the Westin categories, and report that these variables have at best only slightly improved predictive power.

The rest of this paper is organized as follows. In the next section, we provide background information on the Westin Privacy Segmentation Index, as well as other related work. Next, we describe the methodology for our survey as well as describing supplementary data we gathered, and then we turn to findings. We next explore several potential improvements of the Westin Privacy Segmentation Index. We then discuss the implications of our work and conclude.

2. BACKGROUND

2.1 The Westin Privacy Segmentation Index

Beginning in the late 1970's, Westin conducted numerous privacy-related surveys, refining questions and category definitions over time [25]. In 1995, he introduced the Westin Privacy Segmentation Index (subsequently also called the Core Privacy Orientation Index), which he used for nearly a decade in order to make longitudinal comparisons [24, 25]. Note that the questions are specifically related to a consumer perspective, although they have been widely adopted in broader contexts, e.g., [13, 23]. This culminating set of questions is perhaps the most commonly known and used form of his survey instruments, and it is the one we have chosen to include in our study.

A survey using this index asks participants, "For each of the following statements, how strongly do you agree or disagree?" [1 = Strongly Disagree, 2 = Somewhat Disagree, 3 = Somewhat Agree, 4 = Strongly Agree]:

Q1: Consumers have lost all control over how personal information is collected and used by companies.

Q2: Most businesses handle the personal information they collect about consumers in a proper and confidential way.

Q3: Existing laws and organizational practices provide a reasonable level of protection for consumer privacy today.

Based on their responses to these three questions, Westin used the following procedure for dividing participants into three categories [25]. First, responses to the individual questions are classified as follows:

> For Q1, responses of "Strongly Agree" or "Somewhat Agree" are considered privacy-concerned.
>
> For Q2 and Q3, responses of "Strongly Disagree" or "Somewhat Disagree" are considered privacy-concerned.

Next, participants are categorized according to the following rules:

1. Privacy Fundamentalists: Participants who give privacy-concerned responses to all questions;
2. Privacy Unconcerned: Participants who give responses that are *not* privacy-concerned to all questions;
3. Privacy Pragmatists: All other participants (i.e., participants who give a mix of privacy-concerned and not privacy-concerned responses).

In addition to these three questions, Westin drew on other items in his survey instrument to construct a representation of the categories. The essential meaning of these categories remained the same, although specific details varied over the years [25]. The 2002 Harris report provides the following representative descriptions of fundamentalists, pragmatists, and unconcerned [24]:

> **Privacy Fundamentalists**: At the maximum extreme of privacy concern, Privacy Fundamentalists are the most protective of their privacy. These consumers feel companies should not be able to acquire personal information for their organizational needs and think that individuals should be proactive in refusing to provide information. Privacy Fundamentalists also support stronger laws to safeguard an individual's privacy.
>
> **Privacy Pragmatists**: Privacy Pragmatists weigh the potential pros and cons of sharing information, and evaluate the protections that are in place and their trust in the company or organization. After this, they decide whether it makes sense for them to share their personal information.
>
> **Privacy Unconcerned**: These consumers are the least protective of their privacy – they feel that the benefits they may receive from companies after providing information far outweigh the potential abuses of this information. Further, they do not favor expanded regulation to protect privacy.

2.2 The Privacy Paradox

Numerous studies have documented an attitude-behavior dichotomy (also referred to as the Privacy Paradox), in which participants' privacy-related attitudes are seemingly at odds with their actual or intended behavior, e.g., [41, 4, 3]. Spiekermann et al. compared self-reported privacy preferences (as measured with an instrument building on Ackermann et al.'s work [1]) with actual disclosing behavior during an

online shopping episode, finding that participants did not live up to their self-reported privacy preferences [41]. Acquisti and Grossklags studied the relationship between general privacy attitudes and self-reported adoption of privacy preserving strategies and self-reported past release of personal information, and also found supporting evidence for the attitude-behavior dichotomy [4]. While, as noted above, it is unsurprising that general attitudes would not precisely predict context-specific behaviors [15], the dichotomy appears to apply not only to general attitudes and behavior but also to *specific* attitudes and behaviors: Acquisti and Gross demonstrated a gap between the information participants said they cared about protecting online, and what they were showing publicly on Facebook [3].

A number of studies have also documented an attitude-behavior dichotomy specifically for attitudes as established by the Westin categories, showing gaps between the Westin categories and actual behavior [29], the Westin categories and behavioral intentions [6, 20], and the Westin categories and specific attitudes [23]. Malheiros et al. reported that the Westin categories failed to predict disclosure of personal data items in an online setting [29]. Consolvo et al. reported that the Westin categories were not a good predictor of how participants would respond to requests for their location from social relations [13]. Jensen and Potts found inconsistent correlations between the decision to purchase in hypothetical e-commerce scenarios and the Westin categories (as established by an instrument they developed to classify participants into Westin categories) [20]. Further, in an investigation of California residents' attitudes toward law enforcement's access to cell phone location data, King and Hoofnagle found that the attitudes professed among fundamentalists, pragmatists, and the unconcerned did not align with Westin's descriptions of their attitudes [23].

Researchers have previously argued that the disconnect between general privacy attitudes (as measured by the Westin Privacy Segmentation Index or other instruments) and behaviors may be due to a multiplicity of non-mutually exclusive reasons. The reasons include: instruments such as the Westin Privacy Segmentation Index measure general attitudes, while behaviors are context-specific [15]; individuals may perform privacy calculus and make choices that are privacy-suboptimal because they are the most viable or convenient options, even if they are not in accordance with the individuals' privacy preferences [43, 44]; and/or individuals may lack awareness or information about privacy trade-offs, or be subject to various types of decision-making biases [2, 4]. Specific to the Westin Privacy Segmentation Index, King and Hoofnagle have proposed that the Westin categories and their predictive power may be weakening over time [23].

We build on this previous research on the attitude-behavior dichotomy by exploring the relationship between privacy attitudes (as measured by the Westin Privacy Segmentation Index), behavioral intent, and consequences. We believe this is a novel exploration of whether the attitude-behavior dichotomy extends to consequences.

Numerous studies have analyzed privacy concern, and applied diverse instruments for measuring it [34]. In addition to the Westin Privacy Segmentation Index, researchers have proposed other privacy scales, including the Internet Users' Information Privacy Concerns (IUIPC) scale [30] and the Privacy Concern Scale (PCS) [10], both of which contain more specific questions than the Westin Privacy Segmentation Index. Preibusch has observed that scenarios are one of the common ways of measuring privacy concern [34]. In focus group discussions with a small number of participants, Kwasny et al. introduced six brief scenarios relating to surveillance, location tracking, photo sharing, self-disclosure and relationship building, identity theft, and health disclosure [26]. Ackerman et al.'s work is one of the earliest to report the use of scenarios, and we have drawn on their work for inspiration as a representative example of this approach [1], adding outcomes to enable us to explore participants' responses to specific consequences. We believe this type of use of outcomes is novel and allows us to explore issues which have not been previously investigated, such as the relationship between attitudes and consequences as described above. We also believe we have explored a much wider range of scenarios than previously reported.

3. METHODOLOGY

We ran a two-phase study on Amazon's Mechanical Turk in January and February of 2014, which yielded complete data from 884 participants. We also conducted supplementary surveys on Google Consumer Surveys (GCS). In this section we provide details on our study goals, design, and administration, as well as information about the supplementary data and limitations.

3.1 Study Goals

Our study was broadly designed to explore the relationships among generic privacy attitudes (including the Westin Privacy Segmentation Index), responses to hypothetical scenarios, responses to outcomes, personality traits, and demographics. In this paper, we focus on the relationship between the Westin Privacy Segmentation Index and responses to hypothetical scenarios and outcomes.

Our interest in responses to hypothetical scenarios is not novel; Section 2.2 highlighted several studies that have used scenarios to capture individuals' context-specific privacy preferences. However, in this study, we test individuals' responses to a broader array of scenarios, covering diverse situations, types of data, and possible behaviors. In addition to that, we examine the relatively less explored connection between Westin categories and individuals' reactions to potential consequences arising from privacy-sensitive scenarios. In doing so, our goal was to examine whether, as we induce participants to consider a set of possible consequences of protecting or disclosing data (be those consequences negative or positive), individuals' generic privacy attitudes become relevant predictors of how an individual will subjectively perceive, or react to, those privacy trade-offs.

3.2 Study Design

We designed a two-phase study, which was reviewed and approved by CMU's IRB. Phase I consisted of a survey that included several measures of general privacy attitudes. We aimed to capture a wide range of concerns about online and/or offline contexts. After reviewing numerous scales, we chose four that best balanced the following criteria: frequency of use by other researchers, appropriateness for current online and offline environments, and differentiation from other scales that we included. Specifically, we included: the Westin Privacy Segmentation Index [24, 25]; the Westin Personal Privacy Question which is a single question "How concerned are you about threats to your personal privacy in

America today?" [Very Concerned, Somewhat Concerned, Not Very Concerned, or Not Concerned at All] that was used by Westin several times to measure broad public sentiment (it predates the Privacy Segmentation Index, but Westin continued to use it in at least one study after he introduced the Privacy Segmentation Index) [25]); the Internet Users' Information Privacy Concerns (IUIPC) scale which was introduced in 2004 by Malhotra et al. to measure online privacy concerns [30][1]; and the Privacy Concern Scale (PCS) which was introduced by Buchanan in 2007 to keep up with the changing world of online privacy and asks questions related to common online activities (registration, e-commerce, email) [10][2].

Phase I also included three questions which we designed to measure participants' degree of direct and/or indirect experience with misuse of personal information, drawing on questions such as those reported by Malhotra et al. for inspiration [30].

Finally, Phase I assessed personality characteristics using scales from the psychology literature. After carefully reviewing the literature and numerous personality scales, we chose the nine that best balanced the following criteria: relevance to privacy, prior validation, appropriateness for an online survey, and differentiation from other scales that we included. Specifically, we included: TIPI (Ten Item Personality Inventory) [17]; locus of control [35]; MFT (Moral Foundation Theory) [18]; general disclosiveness (subscales amount, depth and honesty) [19]; generalized self-efficacy [37]; SIRI (Stimulating-Instrumental Risk Inventory) [46]; ambiguity tolerance [28]; hyperbolic discounting [5]; and CRT (Cognitive Reflection Test) [16].

Phase II was administered to the same set of participants, and asked them to imagine themselves in three (out of 20) randomly chosen scenarios (see Appendix A for a complete list). Our focus was to compare general attitudes such as those captured by the Westin Privacy Segmentation Index to specific attitudes and behavioral intention when considering context-dependent scenarios and their respective outcomes. Hence, we created a set of privacy-relevant scenarios that manipulate the type of information participants were asked to imagine divulging or not divulging (financial, health, location, social, or otherwise), the context of the disclosures (for example, the party to whom the information was to be disclosed, online versus offline, whether or not the information was anonymized, when or if the information would be deleted) and the consequences of the disclosure (a range of positive and negative outcomes with different financial, health, social, and other impacts).[3] For example, Scenario 1 entertains the following situation: 'A marketing company offers you $1000 and free genetic testing in exchange for the rights to all your current and future medical records. They will have the right to resell or publish your data (anonymously or with information that could identify you, at their discretion)'.

The main response or dependent variable of this study was the answer to a question about likelihood of disclosure (henceforth **scenario response**). The specific question was "How likely would you be to [perform a given action]?" (on a 5-item Likert scale [1 = Not at all Likely, 2 = Slightly Likely, 3 = Moderately Likely, 4 = Very Likely, 5 = Extremely Likely]). For example, for Scenario 1, the exact wording was "How likely would you be to take the offer?"

We were also interested in additional variables that would allow us to better understand the participants' interpretation of and decision-making regarding the scenarios. Accordingly, in addition to the response variable measuring likelihood of disclosure, we also asked questions about participants' specific feelings about each scenario. Specifically, we asked about their confidence that they could make a good decision; how well they thought they could foresee what might happen if they disclosed the information; how risky they felt it would be to disclose the information; how much choice they felt they had about whether or not to disclose the information; how much control they thought they would have over what happened to the information if they disclosed it; how likely it was that they would be in this situation; and how advantageous/disadvantageous the scenario was overall, in the best case, and in the worst case for themselves, their friends and family, and members of society.

After participants responded to questions about three scenarios, we presented them with three outcomes for each scenario (randomly chosen from sets of scenario-specific outcomes) and asked them to make similar assessments in terms of attitudes and disclosure likelihood as they had originally done for the scenarios alone. For a given outcome, the **outcome response** is the participants' reported likelihood of agreeing to the scenario, assuming this was the only outcome. The outcomes represented a wide range of situations with positive, negative, or neutral implications for privacy or well-being. For example, one of the outcomes for Scenario 1 postulates that, 'Your medical data is combined with that of many others. It is used to find a new cure for a previously deadly disease. Neither you nor anyone in your family has this disease.' The complete text of the scenarios and outcomes appears in Appendix A. Based on cognitive testing in a pilot round, each participant was presented only three scenarios, plus three outcomes for each scenario, in order to minimize learning effects and fatigue. At the end of Phase II, demographic data was collected.

3.3 Survey Administration

We administered the survey on Amazon's Mechanical Turk (MTurk) platform in late January and early February of 2014.[4]

[1]We included three components of this scale, namely control, awareness, and collection. These are the novel components the authors introduced in [30]. The scale contains several additional components which are modifications of previous scales, some of which were originally designed for offline environments; we did not include these because they overlapped with other scales we included and/or because they appear less relevant in the contemporary context.

[2]We included the Privacy Attitudes component of this scale (with slight modifications to align the answer choices with other scales). We did not include the Privacy Behavior component which was less relevant for our purposes because of its focus on the use of specific technical capabilities.

[3]For this exploratory study, we did not manipulate the cross-product of all possible variables, but rather focused on the scenarios and outcomes that are most organic and natural based on a review of media reports, research reports, and our experience with participants' concerns in other studies. Future work would profitably include a more systematic manipulation of such variables.

[4]We also ran a pilot version of the survey in April of 2013, with a nearly identical survey instrument and approximately the same number of participants. We re-ran the survey in

For Phase I, MTurk workers were invited to complete a survey about personality and attitudes for a compensation of $2.50. Workers were required to have the following qualifications: live in the United States, Human Intelligence Task (HIT) approval rate >= 95%, and number of approved HITs >= 100. In the MTurk task description, we did not mention privacy to avoid biasing our population. The average completion time for Phase I was 18 minutes, making the average hourly compensation $8.20. This is roughly on par with the United States minimum wage and consistent with payment standards of the MTurk community. A total of 1000 workers completed the task for Phase 1. After data quality assessment, 27 turkers were removed from consideration due to failing catch questions and/or giving overly uniform answers to a large number of questions in a row. After allowing a week to pass in order to minimize potential priming effects from questions in Phase I, we invited the remaining 973 workers to complete Phase II for a compensation of $3.00. 884 individuals out of 973 (90.85%) recruited for Phase II completed it; data from all 884 of these participants is included in the analysis reported in this paper. The average completion time for Phase II was 17 minutes, making the average hourly compensation $10.66.

Table 1 shows several key self-reported demographic characteristics of this sample.

Table 1: Select demographic characteristics of the survey sample.

Demographic	Category	Frequency
Gender	male	47.07%
	female	40.39%
	other	0.31%
	prefer not to answer	0.21%
	skipped	12.02%
Age	18-24	19.84%
	25-34	38.85%
	35-44	15.01%
	45-54	8.02%
	55-64	5.34%
	65+	0.72%
	prefer not to answer	0.21%
	skipped	12.02%
Education	some HS	0.62%
	HS	9.15%
	some college	31.86%
	college	39.05%
	advanced degree	6.89%
	prefer not to answer	0.31%
	skipped	12.13%
Income in $	<20K	17.16%
	20-45K	29.29%
	45-70K	23.74%
	70-100K	9.56%
	>100K	5.65%
	prefer not to answer	2.57%
	skipped	12.02%

early 2014 (screening by MTurk ID to exclude prior participants) to ensure we had recent data to report, to correct a minor typo in Q3 (we also ran a GCS survey with and without the typo with 1500 participants in each condition and did not find a significant difference), and to test the robustness of the results across multiple administrations of the survey. Results from the pilot were largely similar to those reported in this manuscript, with the minor exceptions noted in Section 4.1, and are not included here for the sake of brevity.

3.4 Supplementary Data

We ran several supplementary studies on Google Consumer Surveys (GCS) to contextualize our analysis. These studies are not core contributions of this work, but are included as useful context for the reader. GCS is a market research tool that supports online surveys [22]. Internet users complete survey questions in order to access premium content, and publishers get paid as their users answer. Answers are anonymous and are not connected to personally identifiable information. Demographics (age, gender or geography) are inferred for some participants; this demographic information can be used to target questions to participants or to weigh the results.

In this paper, we include results from two GCS surveys. For both surveys, we targeted the general population in the United States, and we use raw data rather than weighted data for our analyses, as the inferred demographics may not be accurate [22].

First, we ran a GCS survey with the three questions from the Westin Privacy Segmentation Index with 1,500 participants in January 2014.

Second, we ran a GCS survey with 6,000 participants in February 2014 to explore participants' sensitivity to mentioning specific brands. It contained original and manipulated versions of the three questions from the Westin Privacy Segmentation Index, plus three additional questions about purchasing history and trust. This "Brand Survey" had six conditions (1000 participants per condition). In one condition, participants answered the original Westin questions. In the additional five conditions, participants answered the Westin questions modified to refer to Amazon, PayPal, Safeway, Visa, and Walmart rather than more generic terms such as "companies" or "businesses". After answering the three (modified) Westin questions, participants answered three questions about their frequency of past purchases at the specified company (or "online" for classic Westin), their intent to purchase from the specified company (or "online" for classic Westin) again in the future, and how trustworthy they found the company (consistent with Joinson et al's finding that there is a strong relationship between privacy and trust [21]). The full questions appear in Appendix B.

3.5 Limitations

The quality of responses and the composition of the sample are key issues in survey research. In this paper we focused on US respondents in order to reduce heterogeneity of the sample, and we leveraged MTurk and GCS. MTurk, which has been used in prior usable security and privacy research (e.g., [14, 8]), allowed us to collect data from a large number of diverse participants. Buhrmester et al. found that the MTurk population was significantly more diverse than typical American college samples and that using MTurk could result in data at least as reliable as that obtained using traditional methods [11]. Paolacci et al. similarly found evidence that MTurk yielded data comparable in quality to surveying on a university campus [33].

GCS also has limitations, for example its use of inferred demographics and the context in which questions are asked (brief surveys to access premium content) [22]. Nonetheless, some initial reports about GCS are encouraging. The Pew Research Center compared results for questions on a variety of subjects asked in telephone surveys to those obtained using GCS [22]. The median difference between results ob-

tained from Pew Research surveys and GCS was 3 percentage points, and the mean difference was 6 points. They also reported that the demographic profile of Internet users who respond to GCS is similar to that of Internet users in Pew Research Center surveys, and that technological use profiles are also fairly similar. A white paper from Google reports that GCS performed favorably against both a probability based Internet panel and a non-probability based Internet panel, based on several benchmarks [31]. As another example, New York Times' blogger and statistician Nate Silver reported that out of a wide selection of polls, GCS election polls ranked second in terms of accuracy and lack of bias in predicting the 2012 election results [39]. Further, Schnorf et al. administered a questionnaire with several identical privacy questions to multiple panels and report that the levels of privacy concern for both GCS and MTurk respondents were fairly similar to those of respondents in nationally representative samples [36].

Despite these encouraging findings regarding both MTurk and GCS, neither is likely to comprise a statistically representative sample of the general population. Callegaro et al. argue that not only do univariate statistics often vary across samples, but predictive relationships (including magnitude) can vary as well [12]. Future work would benefit from validation in a representative (or different) population, as well as investigation of cross-cultural issues.

Further, although hypothetical scenarios are often used for measuring privacy concern [34], clearly they do not directly measure behavior or attitudes. It would be valuable to extend our work by testing the predictivity of the Westin Privacy Segmentation Index for other indicators of privacy concern. Finally, as with all negative results, a definitive conclusion can not be drawn; our failure to find a correlation does not mean that none exists.

4. FINDINGS

In this section, we present data on the Westin Privacy Segmentation Index, and responses to scenarios and outcomes.

4.1 Westin Privacy Segmentation Index

Figure 1 shows the distribution of responses to the three Westin questions.[5] In Q1, agreement is privacy-concerned, while in Q2 and Q3, disagreement is privacy-concerned. Taking that into account, all three distributions have the same mode (the second-most concerned bucket).

Figure 2 shows the distribution of the Westin categories. Approximately 49% of participants are fundamentalists, 40% are pragmatists, and 10% are unconcerned.[6][7] For comparison, in Table 2 we include the distribution of Westin categories in several other surveys: the GCS survey we ran with only the Westin Privacy Segmentation Index (GCS1); the condition of the GCS Brand Survey we ran which began with the three unmodified questions from the Privacy Segmentation Index (GCS2); results from Westin's 2003 survey administered by Harris Interactive (we were not able to determine full details of this sample) [25]; and results from Westin's 2001 survey administered by Harris Interactive to 1529 members of the Harris Poll Online database, which were then weighted (although the details of the weighting are not fully provided for proprietary reasons) [24]. We provide these numbers so that the reader may better contextualize our results by making a qualitative comparison, but given the varying compositions of the samples it is difficult to draw any definitive conclusions. It appears to be the case that the MTurk population may contain more fundamentalists than the GCS population and the populations tested by Westin in 2003 and 2001. However, it is unclear whether this higher number is simply due to biases in the MTurk population, or whether it is in fact a more accurate representation of current national sentiment. Investigation with a nationally representative sample would be costly but informative.

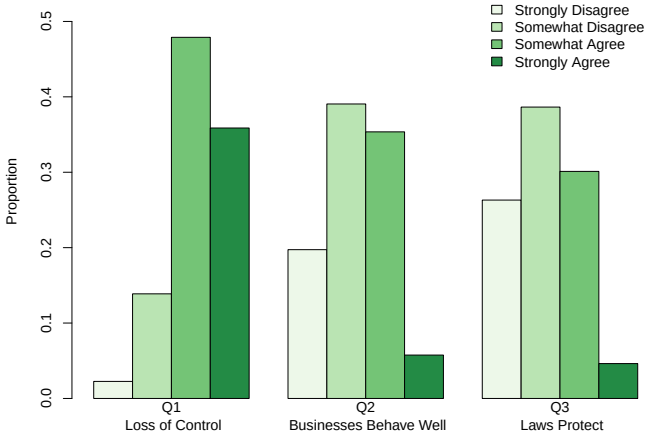

Figure 1: Distribution of raw scores for the three Westin questions. Note the mode of each distribution is the second-most concerned bucket.

We explored whether demographic variables predicted participants' Westin categories or their responses to the individual Westin questions. (Here and throughout, by 'predictive' we mean the ability to accurately predict the previously unobserved value of y based on the value of x given the observed relationship between the two variables in our

[5]For ease of reference we introduce brief précis for the three questions (e.g., 'Loss of Control' for Q1).

[6]Percentages do not sum to 100% due to missing responses.

[7]The alert reader may wonder if Snowden's revelations about NSA surveillance beginning in June 2013 affected the results [27]. Because we had conducted a pilot in April of 2013, we were able to compare data from before and after these events. There are marginally significant differences in responses to Westin's Q1 and Q3 before and after the NSA surveillance revelations. We found that both Q1 and Q3 showed increased concern of about 0.08 on the Likert scale after the NSA surveillance revelations, even after controlling for demographic differences. We did not find significant differences for Q2, nor did we find significant differences for the Westin categories (P-value: 0.8463 for the X^2 test). The minor shift in concern captured by Q1 and Q3 did not appear in the categorization because there was a shift to more extreme positions (from 'Somewhat Agree' to 'Strongly Agree' for Q1 and from 'Somewhat Disagree' to 'Strongly Disagree' for Q3) but not a change in polarity (the distribution of all 'Agree' answers and all 'Disagree' answers for a given question was relatively stable), and the Westin categorization rules rely on polarity. Overall, we found very few differences before and after the NSA surveillance revelations. For example, we found no changes in scenario response, with the exception of Scenario 13 about government surveillance of email, which participants were less likely to support post-revelation.

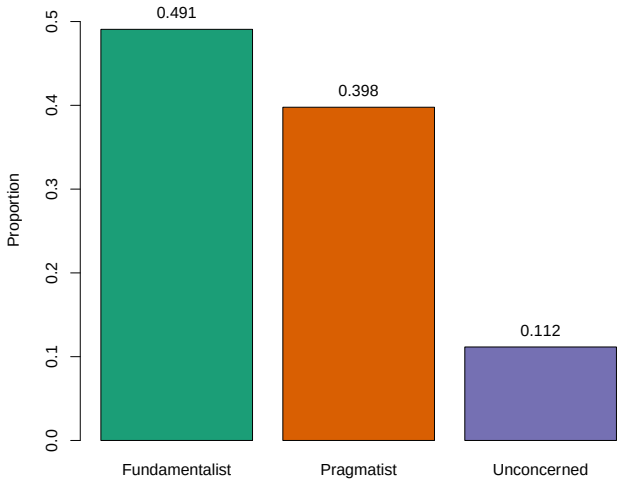

Figure 2: Distribution of the Westin categories.

Table 2: Distribution of the Westin categories in select data sets.

DataSet	Fundamentalist	Pragmatist	Unconcerned
MTurk '14	49%	40%	10%
GCS1 '14	38%	57%	6%
GCS2 '14	37%	58%	5%
Harris-Westin '03	26%	64%	10%
Harris-Westin '01	34%	58%	8%

sample.) Participants self-reported age, gender, education level, income, area where raised, area currently living, employment, religion and ethnicity. These demographic variables do not appear to be correlated with the Westin scale in our sample. No significant demographic predictors were found for any of the three individual Westin questions. We tested the association between the Westin categories and all the demographic variables using a X^2 test [40] with Monte-Carlo p-values because of the small counts in some table cells. Again, no significant associations were found.

We also explored whether any of the personality traits predicted participants' Westin categories or their responses to the individual Westin questions using separate one-way ANOVA models for each trait [40]. Full results are not shown due to space limitations, but in brief we found that purity, in-group, and authority (three dimensions of the MFT scale) have the highest predictive power for the three Westin categories, although the effects are modest (fundamentalists and unconcerned differ by at most 0.4 standard deviation units for any of the traits). The other three variables that show strong evidence of being correlated with the categories are locus of control, emotional stability and CRT; again the effects are modest. These six personality traits differentiated the fundamentalists from the pragmatists and unconcerned, although they revealed little differentiation between the latter two categories. These six traits had p-values < 0.00003 and were significant after correcting for multiple testing using the Bonferroni adjustment (0.05/76) [38]. Regarding the individual questions, similar results to the categories were found for Q2 and Q3, but not for Q1.

The reader will notice that we perform multiple tests for most of our analyses. Given the nature of this large exploratory study, it is critical to test a broad set of pre-defined hypotheses to narrow down the scope of studies that will follow. We are aware of the dangers of data snooping [45] and refrained from running additional analyses to discover 'interesting' results. In all cases, we used a Bonferroni correction to control the Type I error at the nominal level of 0.05 and to avoid an excessive number of false positive findings [38].

4.2 Scenarios

One of the main goals of this study was to examine the relationship between scenario response (i.e., likelihood of disclosure for a given scenario) and Westin categories. Each participant responded to three randomly chosen scenarios presented in a random order. The average sample size per scenario was 128 (min:109 and max:164). No significant differences were observed for any of the 20 scenario responses between Westin categories. Results are summarized in Figure 3. The x-axis lists the 20 scenarios and the y-axis shows scenario responses on a 5-point Likert scale, where 1 indicates 'Not at all Likely' to disclose and 5 indicates 'Extremely Likely' to disclose. Raw responses to scenarios are shown as colored dots (jittered) with three colors corresponding to the three Westin categories. Three solid colored lines trace the means for each category across the 20 scenarios. If Westin categories were significantly correlated with scenario responses, we would expect substantial divergence between the means lines. However, the data supports highly overlapping and crossing means and provides little evidence to the contrary. A formal analysis using one-way ANOVA models to test for differences in means between the three Westin categories for each scenario separately provides further evidence for the lack of association. Several marginally significant differences (Scenarios 3, 7, 11 and 13) disappear after the Bonferroni correction.

Proportions of variance explained by the ANOVA models, R^2's, range from 0% to 7% with a mean of 2% and give another indication of the insufficient ability of Westin categories to predict scenario responses. R^2 is a measure of the goodness of fit and is computed as a ratio of variance (in the response) that is attributed to the Westin categories divided by the total variance in the response.

Distributions of Westin categories within each response category are shown in the right margin of Figure 3. These are shown mainly for qualitative comparison to give the reader a sense of how the sizes of the three Westin categories differ for different response classes after combining all scenarios. If the Westin categories were predictive of the response, we would expect participants who answered 5 (Extremely Likely to disclose) to lean towards the unconcerned category, while those who answered 1 (Not at All Likely to disclose) would be mostly fundamentalists. We do not, however, observe substantial differences in terms of the distribution of Westin categories between these groups of participants.

The three Westin questions are framed in terms of consumer privacy, so we were also interested in comparing the predictive accuracy of the Westin categories for scenarios related to consumer privacy versus scenarios that did not have a consumer aspect. Two of the authors coded our 20 scenarios into three groups, with 100% agreement: three consumer-related scenarios (1, 3 and 4), six marginally consumer-

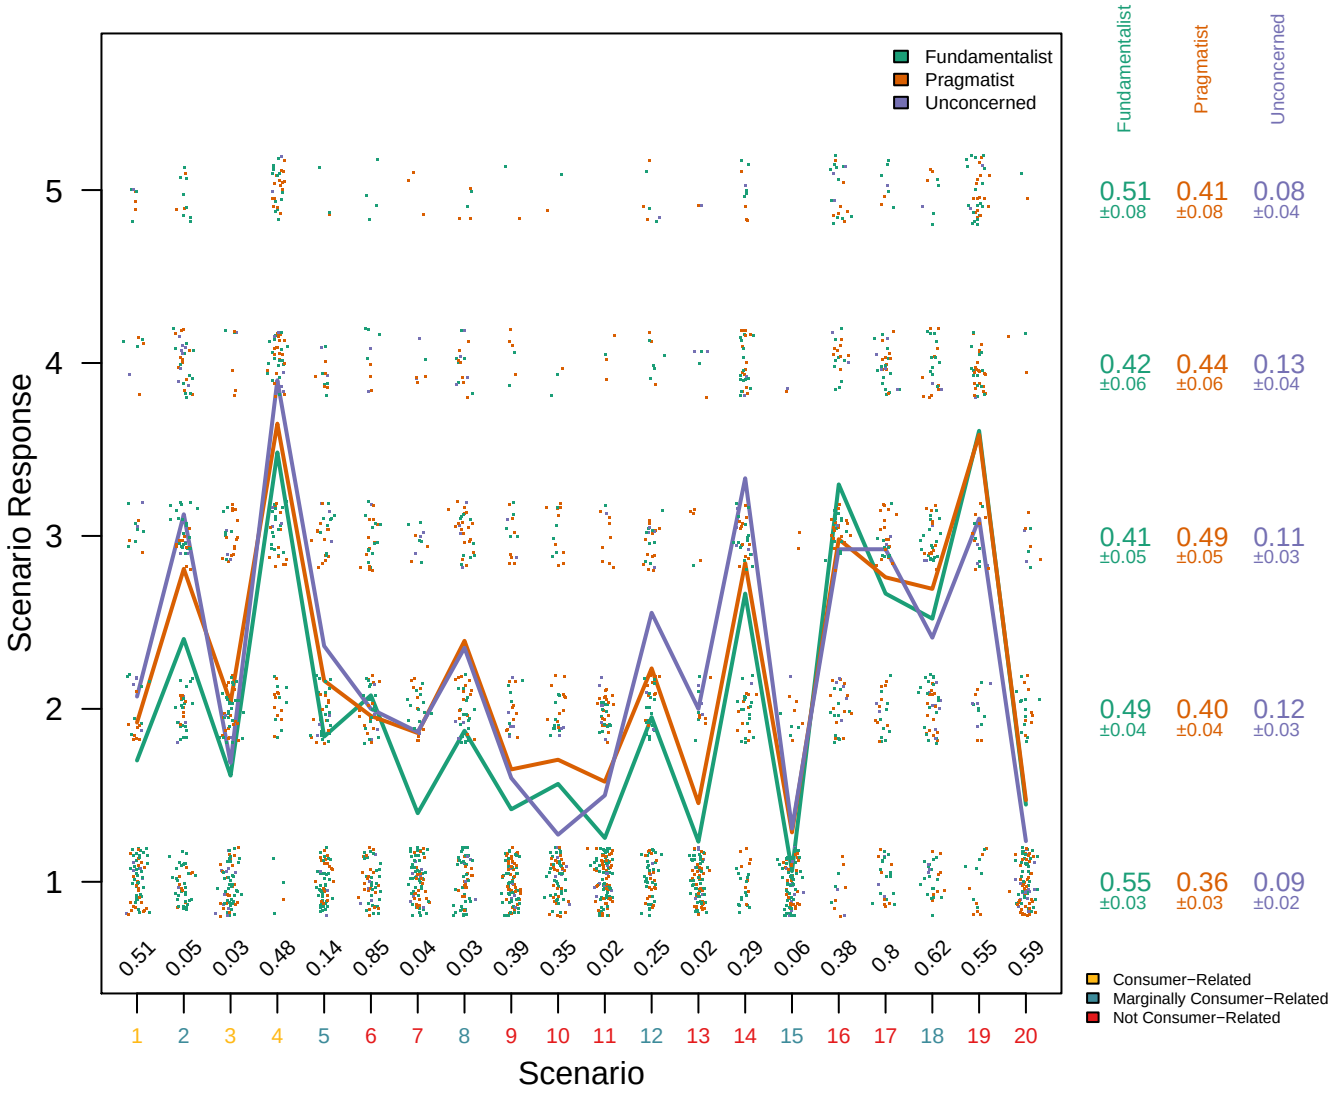

Figure 3: Only small differences (none significant with p-values shown at the bottom) were observed in the scenario response between the three Westin categories. Individual colored dots represent jittered scenario response with colored lines indicating the means for each segment. Scenario numbers at the bottom in different colors indicate the inferred scenario type and show no apparent patterns. In the right margin, the distribution of Westin clusters among each response category is shown with ± two standard deviations.

related scenarios (2, 5, 8, 12, 15 and 18), and 11 non-consumer related scenarios. These three types of consumer-relevance are shown in different colors in the labels for the x-axis of Figure 3. No clear difference emerges in terms of how Westin categories differ by consumer-relevance.

Just as the Westin categories are not predictive of the scenario response, individual Westin questions also show no significant associations (data not included for the sake of brevity). For all 20 scenarios, the proportion of variance explained by the three Westin questions ranges between 1% and 8%. We also note that the Westin categories do not appear to systematically predict any of the 15 scenario variables we collected. Only 6 of the 300 scenario-variable combinations (20 scenarios x 15 variables) had p-values < 0.001, and just 4 were significant after the Bonferroni adjustment.

4.3 Outcomes

In order to understand how increasingly specific information about situations affects responses, three randomly selected outcomes for each scenario were presented to participants. In total, 74 outcomes (3-5 per scenario) were considered and the average responses for each outcome for each Westin category are shown in Figure 4. The outcome response variable is the response to the question "How likely would you be to [disclosure specifics varied by scenario], knowing that this would be the only outcome?"

Clusters of three colored bars (one for each Westin category) represent an outcome. Scenario 1, for example, had

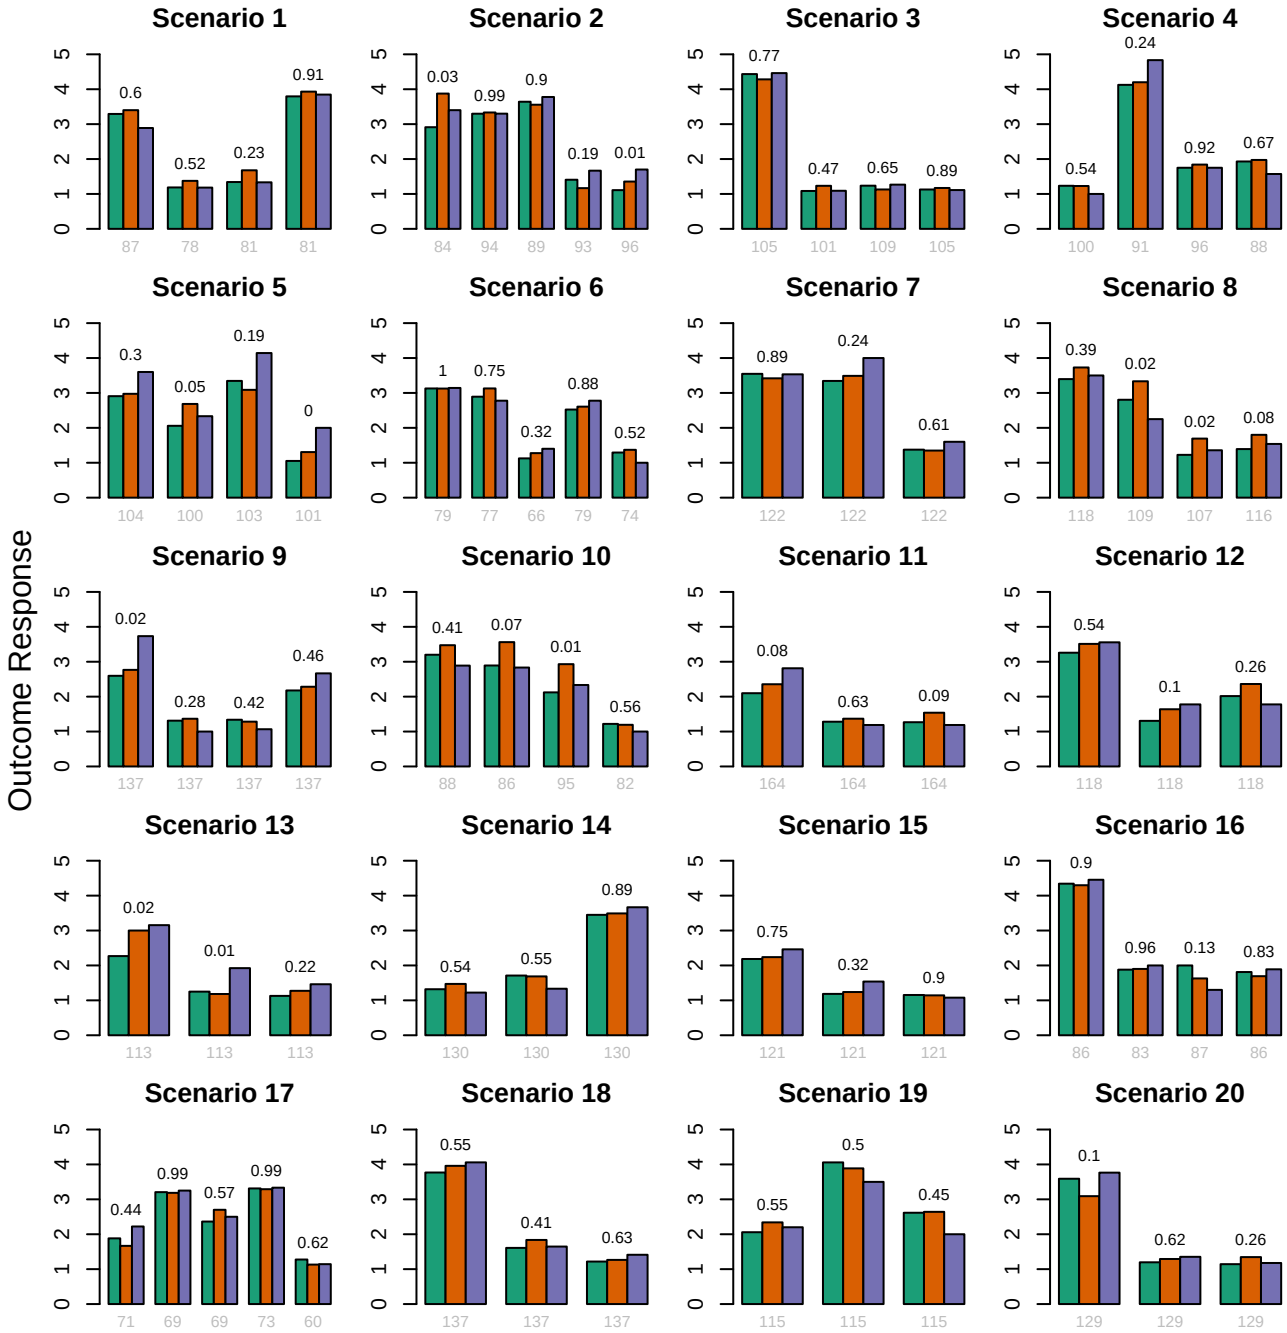

Figure 4: Westin categories by outcomes within each scenario are not significantly different for any of the 74 outcomes. P-values are shown at the top of each cluster of three bars, representing fundamentalists, pragmatists and unconcerned with the same colors as before.

four outcomes, while Scenario 2 had five outcomes. Counts in grey color under each combination of bars show the sample size of each outcome. If Westin's categories were significantly associated with outcome responses, we would observe bars of different colors having significantly different heights, but, as the figure shows, there is no systematic difference across the various outcomes of each scenario. P-values from a one-way ANOVA model [40] are shown at the top of each outcome cluster and indicate how different the Westin categories are in their response. Most outcomes do not show significant differences between the categories and none are significant after the Bonferroni correction. Please keep in mind again that with 74 tests, we would expect just under four of them to be significant prior to the Bonferroni correction even without any true differences. Overall, our results support the conclusion that Westin categories do not capture much of the variation present in the outcome response. As with scenarios, no significant differences were found for

the consumer-relevance of the outcome (p-value 0.5182).

Furthermore, individual Westin questions do not show significant associations with the outcome responses either. The total proportion of variance explained by the three individual Westin questions collectively ranges between 0.2% (Scenario 16, Outcome d) and 15% (Scenario 5, Outcome d) with a mean of 4.2%. Finally, although we do not present detailed data due to page limits, we note that the Westin categories also do not appear to predict any of the 5 outcome variables we collected. Only 5 out of 370 outcome-variable pairs (20 scenarios x 15 variables) had p-values < 0.001 and none were significant after the Bonferroni adjustment.

5. CAN THE WESTIN PRIVACY SEGMENTATION BE IMPROVED?

In the previous section, we failed to show a connection between the Westin Privacy Segmentation Index and responses to hypothetical scenarios and outcomes. In addition to the explanations that have been previously raised, in this section we explore three other possibilities. First, we explore whether different segmentation rules might yield a segmentation that is more predictive of responses to our hypothetical scenarios and outcomes. Second, we explore whether slightly modified versions of the Westin questions (made more specific by providing names of actual companies) yield different responses than the original questions. Third, we explore whether any of the other variables we measured were more predictive than the Westin Privacy Segmentation Index.

5.1 Data-Driven Segmentation

The Westin categories did not capture a significant amount of variation for responses to either scenarios or outcomes. However, it is possible that the three individual Westin questions capture more predictive information about individuals' privacy concerns but the segmentation rules themselves are not optimal and lead to an inferior separation ability.

To investigate this issue, we carried out a clustering of the Westin data using the k-means clustering algorithm with three clusters (other researchers have also used k-means clustering to classify subjects according to their privacy attitudes, e.g., [4, 41]). To visualize the relationship between how participants answer Westin questions and which group they are assigned to by the clustering algorithm, we present Figure 5. The three Westin questions are shown both in rows and columns. For example, the second panel in row 1 corresponds to Q2 in the x-axis and Q1 on the y-axis and shows the joint distribution of responses to these two questions. We invert the responses to Q2 and Q3 so that higher scores indicate more privacy concern. Each dot represents a pairwise response from a single participant, colored according to cluster. Responses are jittered to minimize overlap.

The three clusters found by the algorithm separate quite well, with clear clusters in the bottom left (least concerned, colored green), the middle (moderately concerned, colored blue), and the top right (most concerned, colored red) of each panel. The data-driven segmentation is somewhat different from the Westin one, and the distribution is different as well (cluster sizes are shown below the figure). Table 3 shows what happened to the original Westin categories during the new segmentation. The unconcerned group remains intact in Cluster 3 and receives an additional 52 participants from the pragmatist category. The pragmatist category loses

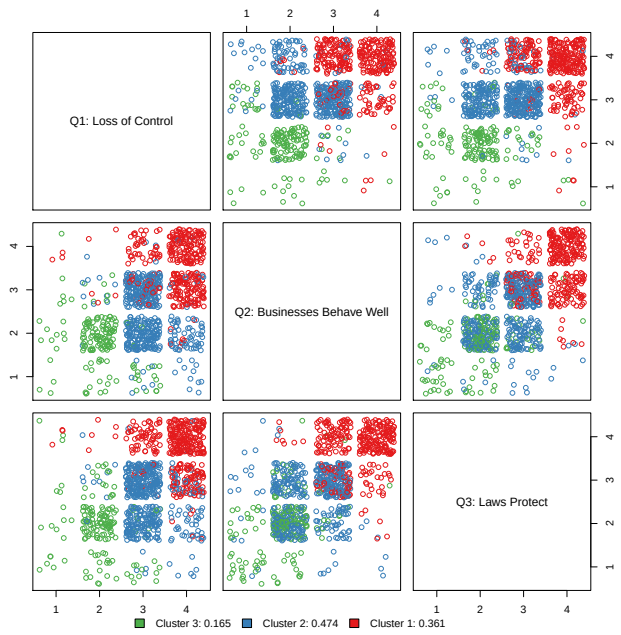

Figure 5: k-means clustering of Westin data with responses to Q2 and Q3 inverted so that higher scores on all questions indicate elevated concern. The x-axis and y-axis show the 1-4 Likert scale. Clear pairwise separation between clusters can be seen. Cluster sizes are shown below the figure.

Table 3: Data-driven segmentation in columns versus Westin categories in rows. A large portion of the difference is the split of the original fundamentalist category into Clusters 1 and 2.

	Cluster 1	Cluster 2	Cluster 3
Fundamentalist	329	146	0
Pragmatist	20	313	52
Unconcerned	0	0	108

an additional 20 participants to Cluster 1 (the new fundamentalist cluster). The largest difference between the two segmentations is the split of the original Westin fundamentalist category into two groups, which reduces the size of the new fundamentalist cluster significantly.

Because Westin prescribed specific segmentation rules, it is interesting to see what rules can be learned from the new segmentation. We use recursive partitioning [9] to that end (Figure 6). Q3 is the most informative of the three questions and is the first condition at the root of the tree. Thus, Q3 is the single variable that best splits the data into the two most homogeneous groups by maximizing the sum of the Gini index for the two nodes. The Gini index measures the impurity of the node in the tree and is defined as

$$1 - \sum p_i^2,$$

where p_i's are proportions of each class (in our case, proportions of each Westin category) in the node. After the initial split on Q3, the split on Q1 is critical for determining the

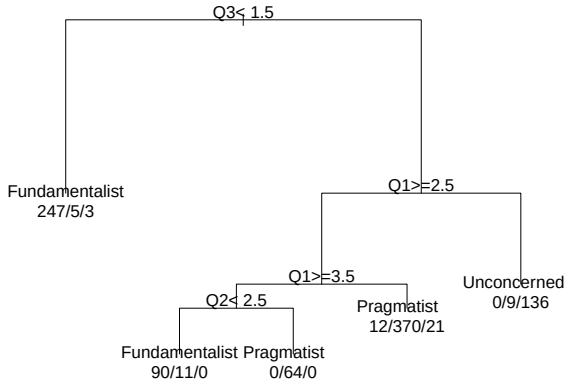

Figure 6: Data-driven segmentation rules for the three Westin questions. True conditions branch to the left and false to the right. The rules differ significantly from Westin's, with Q3 being the most important question to differentiate fundamentalists from others.

unconcerned, while Q2 picks up the remaining differences between fundamentalists and pragmatists. The learned rules for the new segmentation are as follows:

1. Privacy Fundamentalist: 'Strongly Disagree' on Q3 OR ('Strongly Agree' on Q1 and 'Strongly Disagree' or 'Somewhat Disagree' on Q2);

2. Privacy Unconcerned: Q3 is not 'Strongly Disagree' AND ('Strongly Disagree' or 'Somewhat Disagree' on Q1);

3. Privacy Pragmatist: All other participants.

Note that the learned rules do not perfectly reflect the new segmentation. The counts at the bottom in the format 'x/y/z' show how many participants from each cluster (fundamentalist/pragmatist/unconcerned) were classified into a particular category by that sequence of rules. For example, 5 pragmatists and 3 unconcerned answered 'Strongly Agree' on Q3 and are mistakenly attributed to the fundamentalist cluster.

We investigated how well the new segmentation predicts scenario and outcome responses. Results are practically analogous to the Westin categories, with the clusters showing little ability to differentiate participants' self-reported likelihood of disclosing. The clusters show the largest (yet still modest) separation of responses for Scenario 8, and this difference is statistically significant even after the Bonferroni correction (p-value 0.002). Similarly, Outcome b for Scenario 12 is also statistically significant after correction (p-value 0.0006). Overall, performing data-driven segmentation does not result in significantly better response prediction for either scenarios or outcomes.

5.2 Brand Manipulations

We wanted to investigate whether making the Westin questions more specific had an effect. As described above, we ran a GCS survey with 6000 participants in February 2014 to explore participants' sensitivity to mentioning specific brands (Amazon, PayPal, Safeway, Visa, and Walmart) rather than more generic terms such as "companies" or "businesses".

In fact, this small manipulation had a significant effect. Participants were significantly less concerned about privacy when considering a specific company. Table 4 shows differences by brand when compared to the original Westin questions. Brands are sorted by the largest difference in Q1. Very significant and practical differences appear between the five brands and the original general questions. Amazon, by all accounts, received the best marks, where Walmart and Visa yielded values closest to the original questions. Differences in individual questions translate into differences in Westin category frequencies. For the original Westin questions, we observed 37% fundamentalists, 58% pragmatists and 5% unconcerned. The proportion of fundamentalists was smaller for all brands (18% for Amazon and 34% for Walmart), with the proportion of unconcerned growing in all cases (to 25% for Amazon and 16% for Walmart). The trustworthiness variable had the largest effect on the Westin responses (results not shown), but did not explain away the significant differences between the brands after including it in the regression model (along with the other two measured variables about purchasing behavior).

5.3 What predicts disclosure?

We also explored whether personality traits, demographics, situational characteristics, or other privacy scales predicted either scenario or outcome response more effectively than the Westin Privacy Segmentation Index. To examine these relationships, we implemented a mixed-effect model using the *lme4 R* package [7]. Privacy attitudes, personality traits, demographics, and situational variables (all fixed effects) were regressed onto the scenario response along with two random effects (participant and scenario) to account for natural grouping in the data. Results were, perhaps, less encouraging than we hoped.

Analysis of the general privacy attitudinal scales (including the Westin Privacy Segmentation) revealed only small marginal effects that did not seem robust. However, four situational variables, namely, likelihood of the situation occurring, how advantageous the participant perceived the situation would be for them personally, how risky the situation was perceived to be, and how well the participant felt they could foresee the consequences of disclosing had the largest effects on the response. These effects and their corresponding 95% confidence intervals are 0.17 [0.124, 0.2], 0.15 [0.11, 0.195], -0.33 [-0.37, -0.29] and 0.06 [0.02, 0.1] on a Likert scale, respectively. Among the personality characteristics, only disclosure depth (effect size: 0.05 [0.002, 0.1]), disclosure amount (-0.074 [-0.13, -0.02]) and extraversion (0.05 [0.02, 0.09]) were statistically significant from 0. Here, effect size indicates by how much the response changes when the corresponding trait or characteristic changes by one unit. For example, considering the variable for the likelihood of the situation occurring, we would expect participants who answered 'Very Likely' to have, on average, 0.17 higher response scores than those who answered 'Somewhat Likely' given that every other variable remains fixed. The mixed model explains 59% of the response variance using the pseudo-R^2 measure developed as an analogue to the regular linear model. Of this 59%, 38% is attributable to fixed effects (variables we measured), and 21% is attributable to random effects (the participant and the scenario). No multiple testing correction was done here.

We performed the same analysis for the outcome response

Table 4: Differences in mean response for modified Westin questions by brand, as compared to mean response for the original Westin questions. The ± symbol indicates two standard deviations. Adjusted for age and gender.

	Visa	Walmart	Safeway	PayPal	Amazon
Q1: Loss of Control	-0.21 ± 0.081	-0.30 ± 0.080	-0.35 ± 0.081	-0.44 ± 0.080	-0.50 ± 0.080
Q2: Businesses Behave Well	0.12 ± 0.075	0.00 ± 0.075	0.07 ± 0.075	0.26 ± 0.075	0.33 ± 0.075
Q3: Laws Protect	0.20 ± 0.078	0.11 ± 0.078	0.18 ± 0.078	0.32 ± 0.078	0.40 ± 0.078

and obtained very similar results. The outcome mixed model included five additional outcome-specific variables and also the scenario-nested random outcome effect. This model explained about 61% of variance in the outcome response according to the pseudo-R^2 statistic. Of this 61%, 49% is attributable to fixed effects (variables we measured), and 12% is attributable to random effects (the participant and the outcome). Again, how risky the situation was perceived to be (effect: -0.1 [-0.13, -0.08]), how advantageous the participant perceived the situation would be for them (0.03 [0.003, 0.06]) and the likelihood of the situation occurring (0.05 [0.02, 0.08]) were significant scenario effects. All five outcome variables were also significant: how advantageous the participant perceived the situation would be for them personally (0.32 [0.29, 0.35]), how advantageous the participant perceived the situation would be for their friends and family (0.044 [0.01, 0.07]), how advantageous the participant perceived the situation would be for members of society (0.03 [0.006, 0.05]), the likelihood of the outcome occurring (0.18 [0.16, .2]) and how similar an outcome the participant imagined prior to viewing the outcomes (0.037 [0.02, 0.05]).

6. DISCUSSION

The Westin Privacy Segmentation Index is well-established, easy to administer, and yields design-relevant categories. However, consistent with but distinct from previous results, we failed to demonstrate a correlation between the Westin categories and either behavioral intentions or responses to consequences. While our failure to establish a correlation does not mean none exists, certainly the results are not encouraging. At this time, we can not recommend the use of the Westin categories to predict behavioral intentions or responses to consequences. Further, it may be wise to proceed with caution when deploying and interpreting results from the Westin Privacy Segmentation Index for other purposes, unless it has been established to be effective for them. Future work might productively explore whether alternative (e.g., [30, 10] or novel instruments (particularly those considering context [32]) have greater predictive power for both behavioral intentions and consequences.

While the lack of predictive power of Westin's categories across the hypothetical scenarios we presented to our participants is consistent with previous evidence of a gap between attitudes and behavioral intentions, our results also suggest a previously unreported dichotomy between attitudes and consequences. This lack of predictive power relative to actual outcomes can be interpreted in at least two different (and perhaps opposing) manners, suggesting the need for further research. One interpretation suggests that individuals reactions' are based on context-sensitive cost-benefit analyses (encompassing and mediated by complex factors such as systemic biases in decision-making) that are not captured by generic broad privacy attitudes. Another interpretation suggests that the Westin categories may instead capture some underlying, subjective, and deep-seated preferences for privacy that go beyond the so-called privacy calculus, and which may not be fully accounted for by the actual pros and cons of protecting or revealing data. We intend to investigate this further in future research.

A possible implication of these combined findings is that privacy segmentations, or privacy "personas," may inherently face ceilings in terms of their ability to predict privacy choices across diverse real life privacy conditions: there is an unavoidable trade-off between the clustering of preferences that privacy segmentations attempt to construct, and the specificity and heterogeneity of context-specific decisions. At the same time, said segmentations and personas may nevertheless help capture something deep and relevant about people's view of and preferences about privacy.

Finally, our scenarios and outcomes appear to be useful for studying participants' behavioral intentions and responses to consequences. We hope that this instrument may be useful to other researchers. For example, it might be used to explore the predictive value of novel segmentations, or to investigate whether an attitude-consequence gap appears for other instruments.

7. CONCLUSIONS

Previous research has established an attitude-behavior dichotomy, in which participants' broad privacy attitudes as measured by instruments such as the Westin Privacy Segmentation Index are seemingly at odds with their actual or intended privacy-related behaviors. However the relationship between attitudes as measured by the Westin Privacy Segmentation Index and specific consequences has not previously been explored in the literature. We conducted a survey to explore the relationship between the Westin Privacy Segmentation Index and participants' responses to a wide range of hypothetical scenarios and outcomes. We did not find evidence that either the individual questions or the derived categories of the Westin Privacy Segmentation Index are predictive of either participants' behavioral intent or their reaction to specific consequences, suggestive of both an attitude-behavior dichotomy and an attitude-consequence dichotomy. Future research might productively explore the inherent limitations of instruments for measuring broad privacy attitudes, while at the same time considering whether these attitudes capture underlying preferences that are not fully accounted for by contextual or practical considerations.

8. ACKNOWLEDGMENTS

We are grateful to Eyal Peer, Aaron Sedley, Jessica Staddon, and Joshua Tabak for inspiration and advice.

9. REFERENCES

[1] M. S. Ackerman, L. F. Cranor, and J. Reagle. Privacy in e-commerce: Examining user scenarios and privacy preferences. In *Proceedings of the 1st ACM Conference on Electronic Commerce*, pages 1–8. ACM, 1999.

[2] A. Acquisti. Privacy in electronic commerce and the economics of immediate gratification. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 21–29. ACM, 2004.

[3] A. Acquisti and R. Gross. Imagined communities: Awareness, information sharing, and privacy on the Facebook. In *Privacy Enhancing Technologies*, pages 36–58. Springer, 2006.

[4] A. Acquisti and J. Grossklags. Privacy and rationality in individual decision making. *IEEE Security & Privacy*, 2:24–30, 2005.

[5] N. Ashraf, D. Karlan, and W. Yin. Tying Odysseus to the mast: Evidence from a commitment savings product in the Philippines. *The Quarterly Journal of Economics*, 121(2):635–672, 2006.

[6] N. F. Awad and M. Krishnan. The personalization privacy paradox: An empirical evaluation of information transparency and the willingness to be profiled online for personalization. *MIS quarterly*, 30(1), 2006.

[7] D. Bates, M. Maechler, B. Bolker, and S. Walker. *lme4: Linear mixed-effects models using Eigen and S4*, 2013. R package version 1.0-5.

[8] C. Bravo-Lillo, L. F. Cranor, J. Downs, S. Komanduri, and M. Sleeper. Improving computer security dialogs. In *Human-Computer Interaction–INTERACT 2011*, pages 18–35. Springer, 2011.

[9] L. Breiman, J. Friedman, R. Olshen, and C. Stone. *Classification and Regression Trees*. Wadsworth and Brooks, Monterey, CA, 1984.

[10] T. Buchanan, C. Paine, A. N. Joinson, and U.-D. Reips. Development of measures of online privacy concern and protection for use on the Internet. *JASIST*, 58(2):157–165, 2007.

[11] M. Buhrmester, T. Kwang, and S. D. Gosling. Amazon's Mechanical Turk a new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, 6(1):3–5, 2011.

[12] M. Callegaro, R. Baker, J. Bethlehem, A. S. Goritz, J. A. Krosnick, and P. J. Lavrakas, editors. *Online Panel Research: A Data Quality Perspective*. Wiley, 2014.

[13] S. Consolvo, I. E. Smith, T. Matthews, A. LaMarca, J. Tabert, and P. Powledge. Location disclosure to social relations: why, when, & what people want to share. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 81–90. ACM, 2005.

[14] S. Egelman, A. Sotirakopoulos, I. Muslukhov, K. Beznosov, and C. Herley. Does my password go up to eleven?: The impact of password meters on password selection. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, pages 2379–2388. ACM, 2013.

[15] M. Fishbein and I. Ajzen. *Belief, Attitude, Intention and Behavior: An Introduction to Theory and Research*. Addison-Wesley, Reading, MA, 1975.

[16] S. Frederick. Cognitive reflection and decision making. *Journal of Economic Perspectives*, pages 25–42, 2005.

[17] S. D. Gosling, P. J. Rentfrow, and W. B. Swann. A very brief measure of the big-five personality domains. *Journal of Research in Personality*, 37:504–528, 2003.

[18] J. Graham, J. Haidt, and B. A. Nosek. Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, 96:1029–1046, 2009.

[19] B. Grams. Privacy concerns and personality traits influencing online behavior: A structural model. *Dissertation Abstracts International Section A: Humanities and Social Sciences*, 66(7-A):2421, 2006.

[20] C. Jensen, C. Potts, and C. Jensen. Privacy practices of Internet users: Self-reports versus observed behavior. *International Journal of Human-Computer Studies*, 63(1):203–227, 2005.

[21] A. N. Joinson, U.-D. Reips, T. Buchanan, and C. B. P. Schofield. Privacy, trust, and self-disclosure online. *Human–Computer Interaction*, 25(1):1–24, 2010.

[22] S. Keeter and L. Christian. A comparison of results from surveys by the Pew Research Center and Google Consumer Surveys. *The Pew Research Center for the People and the Press*, November 2012.

[23] J. King and C. J. Hoofnagle. A supermajority of Californians support limits on law enforcement access to cell phone location information. *Social Science Research Network*, 2008.

[24] D. Krane, L. Light, and D. Gravitch. Privacy on and off the Internet: What consumers want. *Harris Interactive*, 2002.

[25] P. Kumaraguru and L. F. Cranor. Privacy indexes: A survey of Westin's studies. *ISRI Technical Report*, 2005.

[26] M. Kwasny, K. Caine, W. A. Rogers, and A. D. Fisk. Privacy and technology: Folk definitions and perspectives. In *CHI'08 Extended Abstracts on Human Factors in Computing Systems*, pages 3291–3296. ACM, 2008.

[27] S. Landau. Making sense from Snowden: What's significant in the NSA surveillance revelations. *IEEE Security and Privacy*, 11(4):54âĂŞ–63, 2013.

[28] A. MacDonald. Revised scale for ambiguity tolerance: Reliability and validity. *Psychological Reports*, 26:791–798, 1970.

[29] M. Malheiros, S. Preibusch, and M. Sasse. 'fairly truthful': The impact of perceived effort, fairness, relevance, and sensitivity on personal data disclosure. In *Proceedings of the 6th International Conference on Trust & Trustworthy Computing (TRUST 2013)*, pages 250–266, 2013.

[30] N. K. Malhotra, S. S. Kim, and J. Agarwal. Internet users' information privacy concerns (IUIPC): The construct, the scale, and a causal model. *Information Systems Research*, 15(4):336–355, 2004.

[31] P. McDonald, M. Mohebbi, and B. Slatkin. Comparing Google Consumer Surveys to existing probability and non-probability based Internet surveys. *Google White Paper*.

[32] A. Morton and M. A. Sasse. Privacy is a process, not a PET: A theory for effective privacy practice. In

Proceedings of the 2012 Workshop on New Security Paradigms, pages 87–104. ACM, 2012.

[33] G. Paolacci, J. Chandler, and P. G. Ipeirotis. Running experiments on Amazon Mechanical Turk. *Judgment and Decision making*, 5(5):411–419, 2010.

[34] S. Preibusch. Guide to measuring privacy concern: Review of survey and observational instruments. *International Journal of Human-Computer Studies*, 2013.

[35] J. Rotter. External and internal control. *Psychology Today*, 5(1):37–42, 1971.

[36] S. Schnorf, A. Sedley, M. Ortlieb, and A. Woodruff. A comparison of six sample providers regarding online privacy benchmarks. In *Proceedings of the Workshop on Privacy Personas and Segmentation at SOUPS 2014*, 2014.

[37] R. Schwarzer and M. Jerusalem. Generalized self-efficacy scale. *Measures in Health Psychology: A User's Portfolio. Causal and Control Beliefs*, pages 35–37, 1995.

[38] J. P. Shaffer. Multiple Hypothesis Testing. *Annual Review of Psychology*, 46(1):561–584, 1995.

[39] N. Silver. Which polls fared best (and worst) in the 2012 presidential race. *The New York Times*, November 2012.

[40] M. Sirkin. *Statistics for the Social Sciences*. SAGE Publications, 1995.

[41] S. Spiekermann, J. Grossklags, and B. Berendt. E-privacy in 2nd generation E-commerce: Privacy preferences versus actual behavior. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 38–47. ACM, 2001.

[42] J. Turow. Americans & online privacy: The system is broken. *Annenberg Public Policy Center, University of Pennsylvania*, 2003.

[43] J. Turow, L. Feldman, and K. Meltzer. Open to exploitation: America's shoppers online and offline. *Annenberg Public Policy Center, University of Pennsylvania*, 2005.

[44] J. Urban, C. Hoofnagle, and S. Li. Mobile phones and privacy. *UC Berkeley Public Law Research Paper*, 2012.

[45] S. S. Young and A. Karr. Deming, data and observational studies. *Significance*, 8(3):116–120, 2011.

[46] T. Zaleskiewicz. Beyond risk seeking and risk aversion: Personality and the dual nature of economic risk taking. *European Journal of Personality*, 15(S1):S105–S122, 2001.

APPENDIX

A. SCENARIOS AND OUTCOMES

1. A marketing company offers you $1000 and free genetic testing in exchange for the rights to all your current and future medical records. They will have the right to resell or publish your data (anonymously or with information that could identify you, at their discretion)
   (a) Your medical data is combined with that of many others. It is used to find a new cure for a previously deadly disease. Neither you nor anyone in your family has this disease.
   (b) Your data is published with information that identifies you. You lose a job due to your genetic information, which falsely suggests you may later develop a serious medical condition.
   (c) Your data is used to calculate the probability of certain diseases developing within your family. As a result, some of your relatives (but not you) see an increase of several hundred dollars a year in their health insurance premiums.
   (d) Your test results reveal that you have a serious but treatable disease of which you were previously unaware. You receive treatment just in time to make a full recovery.
2. You join an insurance plan which offers you the option of putting all of your health data in a unified healthcare database. All doctors, hospital staff, and emergency personnel will have access to these records without your needing to give any further permission
   (a) You avoid unnecessary duplicate vaccinations because your current doctor can see that you already received them.
   (b) You no longer have to fill out forms to transfer your medical records from one doctor to another.
   (c) Medical researchers combine your data with that of many other patients. The researchers notice geographic patterns and identify the outbreak of an epidemic much earlier than they would have otherwise. The outbreak, which is located far away from you or anyone you know personally, is contained before it spreads widely.
   (d) Marketers get access to the unified healthcare database and start sending advertising to patients being treated for addiction.
   (e) Your child's doctor looks up your health data and sees that you have been treated for depression. She alerts social services that they should look into whether or not you are caring well enough for your child.
3. Your friend tells you about a company that will give you free, customized investment advice. You go to the website, and to sign up you must provide detailed information about your income, credit history, investments, and investment goals.
   (a) You follow the investment advice and make a huge amount of money. You can quit your current job, retire, and travel the world.
   (b) The company sends you advice that is not helpful at all. They later use your information to commit credit card fraud in your name. They also attempt unsuccessfully to access funds in your bank accounts.
   (c) The company sends you advice that is not helpful at all, and sells your information to several banks. The banks use the information to predict the highest interest rates you personally are likely to pay, and send you targeted credit card and loan offers at precisely these rates. You accept one of the offers and end up paying higher interest than you would have otherwise.
   (d) The company sends you advice that is not helpful at all, and sells your information to several banks.

Based on the information you have provided about your investment goals, the banks conclude you are a poor credit risk and deny you a loan.

4. Your favorite retail store offers you a free loyalty card. You will save an estimated 10% on all store purchases you make when you present the card. To obtain the card, you are required to fill out a form with your name, address, and phone number, which may then be associated with a list of your purchases.
   (a) The retail store sells your data to your health insurance company. Your health insurance company analyzes your purchases, and concludes you have a sedentary lifestyle and an unhealthy diet. They raise your insurance rates.
   (b) You start receiving coupons from the retail store for products you frequently purchase. You end up saving 20% on your store purchases during the year.
   (c) Based on your purchasing patterns, the retail store builds a profile of you and sells it to national marketing companies. You receive tailored offers to which you are susceptible, and end up making some purchasing decisions you would not make normally and that you ultimately regret.
   (d) Your nosy neighbor works at the store. Against the company's rules, they look up the record of all your purchases. They learn that you bought some books about which you are slightly embarrassed. They tease you about the books, although they don't tell anyone else.
5. Your friends are all using a social networking application that lets them publicly share their location online, along with their first names. For example, whenever they arrive at a coffee shop or a bar, they can post that they are currently visiting that place. Your friends ask you to start using the application too, so you can coordinate social activities more easily.
   (a) You post that you are at your neighborhood coffee shop. Unbeknownst to you, a good friend is visiting from out of town. Your friend notices your post and stops by the coffee shop to say hi. You have a great time catching up.
   (b) You start receiving email coupons from the places you've visited, as well as shops near those places.
   (c) The editor at your city's newspaper notices that you go to a lot of performances by cool but obscure bands. They invite you to start writing music reviews for the paper, and you eventually become a minor celebrity.
   (d) A con artist looks up all the locations you have posted. They use the information to strike up a friendship with you, and they ask you for money for an investment opportunity. You invest several thousand dollars, and then you find out the investment opportunity was fraudulent. You feel betrayed and you never get your money back.
6. Your state starts offering a special GPS tag that you can attach to your car. If you have the tag, you can use a special fast lane whenever you go through a toll plaza, and your fare will automatically be charged to your account. Also, state and local agencies will be able to see everywhere you drive so they can manage traffic more effectively.
   (a) Traffic engineers study the GPS data from many users, and greatly improve traffic flow, public transit, and parking in your area.
   (b) Your city uses the GPS data to provide real-time traffic information, which saves you approximately 15 minutes of commute time per day.
   (c) The GPS technology reveals that you are speeding and you get a traffic ticket.
   (d) By using the fast lane at the toll plaza, you save approximately 5 minutes of commute time each day.
   (e) The database with drivers' full names and complete history of locations is hacked and made public. Information about a place you visit for personal reasons is revealed.
7. Your state starts offering a miniature digital monitoring device that can be implanted under a person's skin. The device monitors medical data such as heart activity and body temperature, and it also has GPS tracking to determine your location. The data can be accessed by government agencies and medical personnel in order to assist you or others, but it is not in a publicly available database.
   (a) You have an unexpected allergic reaction that requires immediate medical attention. The device detects the problem and alerts emergency medical personnel. They reach you in just a few minutes, and you make a full recovery.
   (b) You get lost while hiking in a remote area, but because you have the device, you are quickly found by rescue personnel and suffer no ill-effects.
   (c) The government compares GPS data from the devices with locations of crimes, in order to identify suspects. Based on your GPS data, you are wrongly accused of a violent crime and brought in for questioning, although you are quickly released.
8. You discover a free application for your cellphone that collects information about your activity and makes suggestions for improving your health. It automatically collects data on your exercise routes, speed, and duration; it lets you take pictures of food you are eating; it lets you track your sleep habits; and it occasionally asks you how you feel. It analyzes, graphs, and maps the data. It posts the data publicly online, without your name.
   (a) The application points out that you are more active and you feel better when you go to bed before 11pm. You change your habits to go to bed earlier every night. Because of this change, you reach your target weight, you are more productive at work, and you feel happier.
   (b) The application combines your data with that of many others in an anonymous way, and reveals that people feel worse when they go for a walk in your neighborhood. Scientists investigate and conclude that your neighborhood has high levels of pollutants from a local factory. The factory is shut down.
   (c) Someone at your workplace browses the publicly available data for people who live in the area. They

figure out which data is yours, and comment on the fact that you go for a walk every day in the park near their house.

(d) The application starts showing you targeted ads for businesses you pass on your daily walk.

9. Your city government proposes to install an extensive network of surveillance cameras and use face recognition technology to identify and track people as they move around the city.

   (a) The police quickly identify and capture a robber in your neighborhood.

   (b) The police arrest you for being in the proximity of a riot in which you did not participate.

   (c) You are turned away at the entrance to a sporting event because your face is very similar to that of someone who is banned from the stadium.

   (d) Election officials use the face recognition system to identify people who try to vote more than once. Because they eliminate this voter fraud, the candidate you are supporting for a local election wins.

10. The police department in your city proposes to purchase and deploy a fleet of small, low-flying unmanned aircraft that will fly around the city collecting audio and visual data. They explain that they can use this data for purposes such as monitoring city infrastructure or detecting unlawful activity, and they do not plan to make it publicly available.

    (a) The police department uses audio data to detect gunshots in a crime-ridden neighborhood, in which neither you nor anyone you know personally live. Because they are notified quickly when and where gunshots occur, the police are able to catch criminals and provide medical care to victims more efficiently. The crime in the neighborhood decreases quickly and numerous lives are saved.

    (b) During a bad storm, the video helps emergency personnel pinpoint key areas that are flooding. Because of this, they are able to build barricades that successfully protect people and property throughout the city that would otherwise have been injured or damaged.

    (c) The police use audio and video to monitor crowds during a large political protest. They are able to see where large numbers of people are building up and predict where riots are about to break out. They deploy additional security forces to these areas, thereby successfully quelling potential riots before they occur. No one is injured or arrested.

    (d) The police department computers are hacked. The hackers post all audio and video publicly online, including a recording of a very unpleasant fight you had with your significant other. Many of your friends and family see the recording and you feel embarrassed.

11. The political party you support wants to collect information about how individuals feel about various issues and candidates, in order to campaign more effectively. They create a website and ask their supporters, including you, to enter the names of people you know along with any information you have about their political leanings. For example, the website suggests that you enter the political orientation of your neighbors based on campaign signs you see displayed in their yards, and that you enter relevant information you glean from personal discussions with people you know.

    (a) The candidate you support wins the presidential election, in part because helpful volunteers such as yourself enter information about their friends and neighbors.

    (b) Your neighbor finds out you entered information about them. They are angry and cancel plans to have you over for dinner.

    (c) The political party sells the information to marketing companies. These companies use the information for targeted advertising, such as marketing guns to gun supporters and marketing liberal magazines to those who support gay rights.

12. A national newspaper starts publishing an online map that shows all political donations made by individuals. Anyone can search the map by name or address to see which causes an individual donated to, and how much. Many people start using it to look up donations made by people they know. You want to donate money to your favorite political candidate, but many of the people you know aren't aware that you support him.

    (a) Many of your friends see that you donated money, and they are inspired to donate to the candidate you support as well. The candidate you support wins the election, in part because of supporters like yourself and your friends.

    (b) Your boss finds out about your political leanings, and you are passed over for a promotion. You are pretty sure it is because your boss is unsympathetic to your beliefs, but you can't prove it.

    (c) Your next door neighbors find out about your political leanings. You hadn't realized it, but they strongly support the opposing party and they had assumed you did as well. Now every time you see them, they try to change your mind about how you are going to vote. They are polite but extremely annoying.

13. The government is considering passing a law to monitor all domestic email communications for security purposes.

    (a) The government identifies and averts a major terrorist attack.

    (b) The current administration analyzes many individuals' email to determine their political leanings. They use the information to redraw district boundaries and change hours at the polls, in order to give their political party an advantage in the next election.

    (c) Based on your email communications, you are wrongly accused and convicted of a crime you did not commit.

14. You cheated on your significant other, and you feel the need to talk to someone about it. You go out with your best friend, and sit in the corner of the bar. You believe no one can hear you. You consider whether or not to speak.

(a) Someone overhears you. Your significant other comes to know that you cheated on him/her, and breaks up with you.

(b) Your best friend reveals your secret to one of their friends that you are not very close to. However, your significant other never comes to know your secret.

(c) You feel better after discussing your secret with your friend. You recommit yourself to your relationship with your significant other.

15. You hear about a fun new game that you can play on your cell phone, and you think you would enjoy it. You learn that in order to play the game, you must enter the full names and email addresses of twenty of your friends.

(a) The gaming company sends email invitations to your friends to play the game with you. Several of them say yes. You enjoy playing the game occasionally, especially with your friends.

(b) The gaming company sells your friends' names and email addresses to a marketing company. Your friends start receiving annoying spam that appears as though it is from you.

(c) The game is a front for a scam. The gaming company sends email to your friends that looks like it is from you. The email says you are travelling internationally and are in trouble, and asks your friends to wire you money. Several of your friends fall for the scam and lose a total of several hundred dollars.

16. You move into a new home. One evening you overhear a loud fight at your next door neighbors' house. It sounds as though it might escalate into violence. You consider making an anonymous phone call to the police to report it.

(a) The police arrive promptly. You later learn that their presence probably prevented a violent episode, and the aggressor has now moved out of the house.

(b) The police arrive, but they don't find evidence of a problem and they leave. The loud fighting does not resume. However, one of your neighbors guesses you were the one who phoned and the next day they seek you out and tell you they are angry with you. You feel intimidated and you are worried they may retaliate against you in the future.

(c) The police arrive promptly. You later learn that the loud fight was actually the television and there was no problem. Your neighbor laughs it off.

(d) The police arrive promptly. You later learn that the loud fight was actually just a discussion about a football game that was on television. The police give your neighbors a ticket for disturbing the peace, and your neighbors have to pay a large fine.

17. You're at a party, and your friend makes a video recording of you doing a funny dance. They ask your permission to post it on a social networking site, saying they will only share it with mutual friends. You're sure it will make your friends laugh.

(a) Someone you've just started dating sees the video. They decide you are too silly for them and stop returning your calls.

(b) Your friends think it is awesome and compliment you on your moves.

(c) The person who posts the video gets the sharing settings wrong and anyone can see it. It goes viral and eventually appears in the mainstream media. You become a minor celebrity, known for being silly.

(d) One of your friends reshares it with a few people. An old friend from high school finds you because you're in the video. They get in touch with you and you're glad to hear from them.

(e) One of your friends reshares the video. It goes viral and is seen by a prospective employer. You don't get the job you were hoping for, because they think you are too silly to do well at the job.

18. You are planning a family vacation. Your friend recently had a great experience using a house swap website, and they recommend you try it. You go to the website and find a beautiful house in a terrific location in the city that you most want to visit. The owners have good reviews on the website from other people who have swapped houses with them in the past, and they are willing to swap houses with you for free at a time that is convenient for you.

(a) The house you visit is wonderful, and you have a great vacation. The family you swap with leaves your house in perfect condition.

(b) The house you visit is wonderful, and you have a great vacation. However, when you return home you can tell the other family riffled through all your things. Nothing seems to be missing, but you feel uncomfortable.

(c) The house you visit is messy and unpleasant, and you have a mediocre vacation. When you return home, you learn that the other family had a big party at your house and the police were called to break it up. The carpet is stained and several minor items are broken or damaged. You are unable to recoup the costs from the other family.

19. Your city is creating a time capsule that will be opened in 100 years. A photographer goes around town taking photos for the time capsule, and they take a photo of you and your significant other in a passionate embrace. They ask your permission to include the photo in the time capsule.

(a) The photo is included in a brochure describing the time capsule. The brochure is mailed to everyone currently living in your city. Your friends tease you and you are mildly embarrassed.

(b) No one sees the photo during your lifetime. When it is viewed in 100 years, it becomes an iconic image of romance in your time period, and you are immortalized.

(c) No one sees the photo during your lifetime. When it is viewed in 100 years, public displays of affection are frowned upon, and your behavior is considered scandalous.

20. You are searching for a job. You find an advertisement for a job that sounds perfect for you, and you start to complete the application online. When you're almost done, the form asks you to provide your social network

login and password so the human resources department can look at your private posts with friends and family. The form explains this will help the human resources department evaluate your fit with the company's culture.

(a) You are invited to interview, and you get the job. It is indeed a perfect fit for you. You make more money than you ever imagined, and you enjoy your work tremendously.

(b) You are invited to interview, and you get the job. However, you quickly discover the company engages in numerous unethical and illegal practices, and you resign before you get too embroiled in their wrongdoing.

(c) You do not get the job. However, one of the employees in the human resources department finds a private and moderately embarrassing photo of you, and posts it publicly on an Internet site that features such photos.

B. BRAND SURVEY QUESTIONS

In one condition, participants saw the original Westin Privacy Segmentation Index questions for the first three questions. In the other five conditions, participants saw the modified versions as specified below for the first three questions (modifications from the original questions are shown in bold font). Participants in all conditions saw the final three questions.

Consumers have lost all control over how personal information is collected and used by [**Amazon, Paypal, Safeway, Visa, Walmart**].

- Strongly Disagree
- Somewhat Disagree
- Somewhat Agree
- Strongly Agree

[**Amazon, Paypal, Safeway, Visa, Walmart**] handle**s** the personal information **it** collect**s** about consumers in a proper and confidential way.

- Strongly Disagree
- Somewhat Disagree
- Somewhat Agree
- Strongly Agree

Existing laws and organizational practices provide a reasonable level of protection for [**Amazon, Paypal, Safeway, Visa, Walmart**] consumer**s'** privacy today.

- Strongly Disagree
- Somewhat Disagree
- Somewhat Agree
- Strongly Agree

How many times have you made a purchase [from Amazon, with Paypal, from Safeway, with Visa, from Walmart, online] within the past 12 months?

- Never
- 1 time
- 2 - 5 times
- 6 - 10 times
- More than 10 times

How likely is it that you will make a purchase [from Amazon, with Paypal, from Safeway, with Visa, from Walmart, online] within the next 12 months?

- Not at all Likely
- Slightly Likely
- Moderately Likely
- Very Likely
- Extremely Likely

How trustworthy [is/are] [Amazon, Paypal, Safeway, Visa, Walmart, online vendors]?

- Not at all Trustworthy
- Slightly Trustworthy
- Moderately Trustworthy
- Very Trustworthy
- Extremely Trustworthy