Taken Out of Context: An Empirical Analysis of Westin's Privacy Scale

Jennifer King
U.C. Berkeley School of Information
102 South Hall
Berkeley, CA 94720-4600
jenking@ischool.berkeley.edu

ABSTRACT

This paper offers the results of an empirical analysis of Alan Westin's privacy scale using survey data taken from a nationally representative survey of American internet users conducted in July 2009. I focus on two questions: 1) Are there any demographic covariates that can be reliably associated with Westin's three privacy categories?; 2) Can Westin's categories be used to make predictions about privacy attitudes, knowledge, or behavior? I use these findings to suggest that Westin's scale does not provide a reliable framework for analysis or prediction outside of its original context. Privacy researchers in the 21st century should seek new tools for measuring the public's privacy attitudes rather than overloading Westin's scales.

1. INTRODUCTION

For over forty years Professor Alan Westin was responsible for the majority of the privacy polling in the United States. Westin conducted over 120 privacy surveys, focusing primarily on consumer information privacy.[1] His dominance in this area follows the publication of his seminal 1967 work, *Privacy and Freedom*, an early herald of privacy issues wrought by new technologies. Westin noted that throughout his career, there "has been a well-documented transformation in consumer privacy attitudes . . . moving concerns from a modest matter for a minority of consumers in the 1980s to an issue of high intensity expressed by more than three-fourths of American consumers in 2001."[1] Through his polling work, Westin developed a Guttman scale to measure privacy attitudes, dividing the American public into three distinct categories based on their answers to three Likert scale questions.[2] These categories are often cited as a basis for characterizing the American public's attitudes towards information privacy, which in turn are used as a proxy for influencing debates about privacy issues. However, despite the popularity of his work, Westin published little data about what factors influenced his classification scheme beyond the three concepts captured in his scale questions. Furthermore, the scale itself has received little critical attention, possibly due in part to the dearth of published analysis offered by Westin and the high cost of replicating his findings through nationally representative surveys.

In this paper, I conduct an exploratory data analysis focusing on two questions: 1) Are there any demographic covariates that can be reliably associated with Westin's three privacy categories?; 2) Can Westin's categories be used to make predictions about privacy attitudes, knowledge, or behavior? The data are taken from a nationally representative survey of American internet users co-authored by my associates and I and conducted for us by Princeton Survey Research Associates [PRSA] in July 2009. We included Westin's privacy scale questions in the survey as well as questions intended to explore respondents' knowledge of information privacy law and policy, privacy attitudes, and related behaviors. Traditional demographic data was captured and is included in the analysis. Using both binomial and multinomial logistic regression, I explore the significance of demographic and privacy-related covariates on respondents classified into Westin's three categories to explore their predictive power.

2. SURVEY DATA

The data analyzed in this paper were originally part of a larger survey of Americans' opinions about and understanding of a variety of online and offline privacy issues. The survey results were published in a series of two reports: "*How Different are Young Adults from Older Adults When it Comes to Information Privacy Attitudes and Policies?*,"[3] and "*Americans Reject Tailored Advertising and Three Activities That Enable It.*"[4] We cast our population net broadly, including people in our study if they were 18 years or older said yes to one of the following questions: "Do you go on online or use the internet, at least occasionally?" and "Do you send or receive email, at least occasionally?" PSRA conducted telephone interviews with a nationally representative, English-speaking sample of 1,000 American adults living in the continental United States. A combination of landline (n=725) and wireless (n=275) random digit dial (RDD) samples was used to represent all adults in the continental United States who had access to either a landline or cellular telephone. The interviews averaged 20 minutes. Based on a seven callback procedure and using the American Association of Public Opinion research (AAPOR) RR3 method, the overall response rates were a typical 18 percent for the landline sample and 22 percent for the cellular sample. Statistical results were weighted to correct known demographic discrepancies. The margin of sampling error for the complete set of weighted data is ±3.6 percent at the 95 percent confidence level. The margin of error is higher for smaller subgroups within the sample. For more details about the sampling method and a complete list of demographic variables I refer readers to either copy of the reports referenced above.

3. WESTIN'S PRIVACY SCALE

Westin's scale began first with his development of the Harris-Westin Distrust Index. He notes that since 1978, "the driving factors behind privacy attitudes, both in general and in specific consumer areas, to be a combination of two orientations: the individual's level of distrust in institutions and fears of technology abuse."[1] The Distrust Index measures "distrust in government, voting, and business" as well as "fears that technology is almost out of control." Westin says that "the higher the Distrust Score,

the more a respondent will express concern about threats to privacy, believe that consumers have lost all control over uses of their information by business, reject the relevance and propriety of information sought in particular situations, call for legislation to forbid various information practices, etc."[1]

In 1995, Westin introduced his privacy segmentation index in his Harris-Westin surveys, "producing a division that essentially mirrored [three] ideological-interest positions."[5] The three are high, balanced, and limited privacy positions, distributed across a spectrum with roughly 20 percent of the public located at the ends and 60 percent at the center "balanced" position. Westin admits (and survey data affirms) to fluctuations in the distribution over time (see Table 1).

Table 1. Fluctuations in Categorical Distributions, 2009-1999

	Berkeley Survey: 2009	Harris Poll: 2003[1]	Harris Poll: 1999[2]
Privacy Fundamentalists	22%	26%	25%
Privacy Pragmatists	58%	64%	54%
Privacy Unconcerned	20%	10%	22%

Classification into one of the three categories is based upon responses (excluding missing, don't know, or refusals) to each of the following three questions (see Table 2).

Question 1: Consumers have lost all control over how personal information is collected and used by companies.

Question 2: Most businesses handle the personal information they collect about consumers in a proper and confidential way.

Question 3: Existing laws and organizational practices provide a reasonable level of protection for consumer privacy today.

Table 2. Westin's Classification Questions, 2009 Berkeley Survey (reported as percentages)

	Strongly agree	Somewhat Agree	Somewhat Disagree	Strongly disagree	DK/Ref.
Q1	20	47	27	4	2
Q2	5	53	32	6	4
Q3	4	50	34	8	4

Privacy Fundamentalists agree (strongly or somewhat) to (1) and disagree (strongly or somewhat) to (2) and (3). The Privacy Unconcerned disagree (strongly or somewhat) with (1) and agree (strongly or somewhat) with (2) and (3). Privacy Pragmatists are those with any other pattern of responses to the questions. I was able to classify a total of 906 respondents based on valid responses to all three questions.

3.1 Analyzing the Categories

In a 2003 book chapter, Westin offered a "conceptual framework for information privacy analysis" based upon "tracking three settings: The political, the socio-cultural, and the personal)."[5] Political privacy is derived from the political (e.g. democratic versus authoritarian) and legal system in which the society is situated. In the socio-cultural realm, Westin mentions factors such as class, race, power, and social status and their impact on privacy (e.g. "the rich can withdraw from society when they wish; the lower classes cannot"). Individual privacy consists of "one's family life, education, social class, and psychological makeup," from which Westin derived four states of privacy: solitude, intimacy, anonymity, and reserve. Notably, he comments that one's state is privacy is constantly changing, dependent upon context and "changing personal needs."

In the remainder of the chapter, Westin discusses privacy in terms of political privacy, noting the historical and political trends that influence popular opinion and stating nothing about the socio-cultural or the personal, leaving the discussion of those dimensions to other scholars. This is not unusual in Westin's publications; to the best of this author's knowledge, Westin, a political scientist by training, did not explore these dimensions with any of his publicly available empirical data.[3] None of his writings about his privacy scale includes demographic covariates or other factors that are associated with his categories. This may prove not to be an omission by accident but one of intention—it is certainly possible that after decades of survey analysis, Westin could not identify any covariates that he could reliably associate with each category. Or, perhaps he simply found them less interesting or relevant than the influence of the political dimension. In Congressional testimony given in May 2001, Westin admitted to less interest in the demographic factors, noting that "while there are fascinating demographic variations of privacy (e.g., women are 10-20 percentage points more intense on most consumer and internet privacy concerns than men), the most important analysis of public attitudes probably involves how the public divides on consumer privacy issues."[1] Overlooking the possible chauvinistic interpretation of this statement (to this author, gender differences in privacy attitudes is a significant finding worth reporting), the analysis I conduct examines possible significant covariates across the personal, socio-cultural, and political dimensions by regressing general demographics as well as scale variables exploring privacy knowledge, behavior, and attitudes.

4. ANALYSIS

After assigning the 906 respondents to one of Westin's three categories, I created four groups of covariates to use in a series of binomial logistic regressions (using each Westin category as a dummy variable). Each set of covariates and results are discussed in turn. After conducting this analysis, I then attempted a set of multinomial logistic regressions using the Privacy Unconcerned

[1] Taylor, Humphrey. "Most People Are 'Privacy Pragmatists' Who, While Concerned about Privacy, Will Sometimes Trade It Off for Other Benefits." The Harris Poll #17, March 19, 2003.

[2] Ibid.

[3] Westin conducted a number of surveys over the years for private sponsors; while he occasionally made references to them in his work, the contents were not published.

as the reference category. All logistic regressions were run in STATA using the SVY prefix to apply frequency weights.

4.1 Independent Variables

I divided the independent variables into the following groups: demographic variables, privacy knowledge, privacy behaviors, and privacy attitudes. Detailed findings from each regression are summarized in the appendix.

The demographic variables were typical for most national opinion surveys, including gender, age, political views (from conservative to liberal), religiosity as measured by frequency of attendance at religious services, education (simplified; no specific options for vocational or graduate education), race (the survey included separate questions to measure Hispanics as well as other groups; I recoded to simply capture White versus non-White), and household income (simplified to above or below $50K/year).

The three groups of privacy covariates are based upon aggregations of questions from the survey that I created post-hoc for this analysis. In the privacy knowledge set, the online and offline questions were a set of true/false statement measuring knowledge about existing privacy laws and policies.[4] The internet skill measure is based on the response to the question "How would you describe your abilities to go online or navigate the internet?" Privacy behaviors included answers (recoded as yes/no, but originally with multiple options) regarding erasing browser cookies, reading online privacy policies, refusing to provide personal information to a business or company because the respondent felt it was too personal, and whether or not the respondent belonged to a social networking website (SNS). Privacy attitudes were measured by responses to five sets of questions: two prescriptive questions asking respondents first whether they thought there should be laws giving people the right to know everything a website knows about them, and second if there should be a law requiring websites and advertising companies to delete all stored information about an individual; a composite question of five attitudinal questions about online targeted advertising;[5] a question asking respondents about their internet service provider sharing one's internet usage details with advertisers; and a question asking respondents how long advertisers should store information about one's internet activity.

4.1.1 Binomial Logistic Regression

I ran four sets of binomial logistic regressions, one each for the demographic covariates, privacy knowledge, behavior, and attitude groups. Please see Tables 7-10 in the Appendix for details. I did not attempt to interact any covariates as I had no theoretical basis for doing so.

Controlling for all of the demographic covariates, only age was significant for two groups: Fundamentalists and the Unconcerned. Increases in age decrease the odds of being an Unconcerned by 3.1%, while each increase in age increases the odds of being a Fundamentalist by 2%. For the knowledge covariates, the only significant coefficient was the online privacy knowledge measure for Fundamentalists; being a Fundamentalist increases the probability of having a higher score on this measure by 14%. Examining privacy behaviors, the only significant covariate was social networking site membership for the Unconcerned; being classified as Unconcerned increased the probability of being a SNS member by 110%. For privacy attitudes, significant covariates were the "right to know" question[6] for both Fundamentalists and the Unconcerned and the ad storage question[7] for Fundamentalists and Pragmatists. Being a Fundamentalist increased the probability of answering yes to the right to know question by 73%, while being Unconcerned decreased it by 47%, and being a Fundamentalist increased the probability of taking a pro-privacy position on the ad storage question by 44%, while being a Pragmatist decreased it by 22%.

4.1.2 Multinomial Logistic Regression

Because membership in each of Westin's three categories is independent, comparing the three groups against one another instead of binomially (one group against the other two combined) in theory should yield more accurate results. I ran a set of multinomial logistic regressions with the same set of covariates, using the Unconcerned as the base category (assuming that changes towards more privacy protective behaviors should yield increased odds as we examine Pragmatists and Fundamentalists). Please see Tables 11-14 in the Appendix for details.

Comparing the output of these regressions against the binomial set yielded interesting results. First, while one might expect that Fundamentalists would consistently show a significant β_0 value when compared to the Unconcerned, considering they are on opposite sides of the privacy spectrum, this was not the case. Fundamentalists' β_0 when comparing with the Unconcerned was only significant in the demographic regression. This result appears to be entirely driven by age. In this regression, the coefficients for age are significant for both Fundamentalists and Pragmatists, whereas in the binomial regressions this was true for only Fundamentalists and the Unconcerned. After examining the distributions of age by category, in this sample the Unconcerned are generally younger than the other two categories (see Table 15 for a summary). It appears the independent analysis teases out this difference.

Counter-intuitively, β_0 values for Pragmatists, the middle value on the privacy spectrum is significant with the knowledge and behavior regressions (but neither β_0 is significant in the attitude regression compared to the Unconcerned). The β_0 values for Fundamentalists are not significant in any of these three regressions. Examining these results in detail, it isn't clear which covariate is driving this difference in the knowledge regression. The significant value in the behavior regression appears attributable to SNS membership.

5. DISCUSSION

While the analysis presented here is by no means conclusive and would benefit from additional scrutiny, I suggest these findings

[4] The alpha values for the composite scales are .65 for the online questions and .69 for the offline questions.

[5] The alpha value for this composite scale is .83. Please see [3] for the questions used to create this scale.

[6] Q: "Do you think there should be a law that gives people the right to know everything that a website knows about them, or do you feel such a law is not necessary?" (Responses: Yes, No, Don't know/refused)

[7] Q: "Do you think there should be a law that requires websites and advertising companies to delete all stored information about an individual, if requested to do so?" (Responses: Yes, No, Don't know/refused)

raise enough questions to call for additional exploration into the robustness and predictive capabilities of Westin's categories. There are limitations with this analysis; for example, the scales I developed (with the exception of the demographic covariates) were not intentionally designed for post hoc analysis (the alphas range from acceptable (<.7) to good (>.8)), and may not be robust measures of privacy knowledge, attitudes, or behaviors. A strength of this analysis is that the dataset is similar to what Westin used to establish his categories: a large, U.S. based nationally representative telephonic survey that adheres to best practices for sampling. Future analyses should rely on similarly sized and sampled datasets in order to maintain consistency with Westin's original sampling frame.

However, despite these limitations, it is surprising that his categories were not reliably associated with any of my three privacy measures. In the original survey, responses to several of the questions used to construct the scales skewed quite high in favor of privacy, in some cases, over 90% answering similarly. One might expect that at least the Fundamentalists might show a consistent association with these measures. One explanation might be that Westin's categories attempt to measure attitude formation at a fundamental level, whereas questions that probe specific outcomes aren't drawing upon the same constructs. This begs the question: how helpful is it to categorize someone as a Privacy Pragmatists if that categorization can't be used to predict where that person's opinion falls on the issues of the day? This is an area where additional research would be helpful for assessing this outcome more broadly.

The overall lack of influence of demographic covariates is notable. As I mentioned earlier, Westin never reported any influence of demographics on his categories. Does this mean that privacy is a value that transcends common demographic divisions? One would hope so—and perhaps this provides evidence that Westin's measurement taps a fundamental attitude measure that is independent of gender, age, race, and other effects. Generally, this is an area that seems ripe for further analysis, as it surprising (but not implausible) that privacy attitudes are fully independent of any demographic factors.

Interestingly, the only covariate that yielded significant results in both sets of regressions (though weakened in the binomial regression results) was age: age followed a linear trend, as Privacy Fundamentalists were older, and the Unconcerned were younger. However, as we found in [4], there was a broad support for pro-privacy positions across all age cohorts, despite a public perception that younger people don't care about privacy. I suspect this disparity, as well as the overall weakness in the predictive capacities of his scale, may be related to both the context and the construction of Westin's questions. The questions were created to explore consumer privacy attitudes prior to the introduction of the internet, capturing the fear that many had of the *existence* of databases in the 1970s and 1980s (Q1: "*Consumers have lost all control over how personal information is collected and used by companies*"). This concern may continue to resonate with older respondents, but the mere fact of data collection (and existence of databases) likely does not inspire the same fear among younger ones. Younger respondents (*i.e.*, those under 50) are likely more sophisticated about the many ways in which they might exert control over their personal information: credit reports; performing self-searches on search engines; using privacy controls or controlling their disclosure on social media. Younger respondents may not agree strongly that they have lost "all" control.

Next, while Westin attempted to track the "political, the socio-cultural, and the personal" elements of privacy concern with his scale, it isn't evident that his questions map to current concerns in these areas. Both Question 2 ("*Most businesses handle the personal information they collect about consumers in a proper and confidential way*") and Question 3 ("*Existing laws and organizational practices provide a reasonable level of protection for consumer privacy today*") rely on an understanding of business practices and legal consequences that is potentially out of scope for many Americans today, but particularly younger ones. In [4], my colleagues and I asked a series of true/false questions probing both on and offline privacy knowledge; regarding online knowledge, only 75 percent of the entire survey sample answered only two or fewer questions correctly, with the youngest adults (aged 18-24) performing far worse than other age cohorts: 88 percent answered only two or fewer correctly. Our offline questions displayed a nearly identical dearth of knowledge: 74 percent of the sample answered two or fewer correctly, with 88 percent of the youngest adults again only answering two or fewer correctly. [4] Sadly, ignorance in this area isn't limited to the young, though younger respondents do perform worse on these measures. The demographics of the American public have changed dramatically since Westin first introduced these questions, and the assumptions he made about the public's general understanding of these concepts needs to be revisited.

Finally, none of these analyses control for a factor that may have an impact on how respondents answer these questions: a prior adverse experience with privacy. Trepte *et al*. [6] found in a longitudinal study of social networking sites that prior negative experiences positively influenced participants' risk assessment and informational privacy behaviors. Cheshire *et al.* [7] found a strong association with having an adverse experience[8] and a negative perception of one's control over their online information. In an attempt to probe the influence of adverse privacy experiences, in a 2010 survey I performed via a Facebook Platform app of Facebook users I asked respondents several questions tracking a variety of negative privacy experiences on social networking sites and created a scale measure for the responses.[8] After performing a series of multiple regressions to control for a variety of covariates, this measure of adverse events was the only factor that reliably allowed us to predict respondents' privacy attitudes. According to the National Criminal Justice Reference Service, in 2012 7 percent of Americans over the age of 16 were victims of identity theft. [9] If one considers the variety of negative privacy events people may experience today online, the impact on a representative national survey may be significant if adverse privacy events in turn produce respondents with stronger pro-privacy concerns or attitudes. It is possible that Westin's questions capture this effect implicitly, albeit accidentally, but an explicit measure would give privacy researchers more insight into the factors that contribute to attitude formation.

To be fair, Westin does not make global, sweeping claims that these categories explain privacy attitudes outside of the consumer information privacy niche. Furthermore, he is explicit (as quoted earlier) in noting his belief that privacy concerns emanate from distrust in institutions and fears of technology abuse. While his

[8] An adverse event included negative experiences with both security and privacy.

categories enable a tantalizingly simple classification of the public into three distinct and tidy categories, we should ask: do they adequately explain the core beliefs that shape individuals' online information privacy attitudes? Or does his scale measure little more than one narrow—and now increasingly anachronistic—view of information privacy?

6. CONCLUSION

This analysis is but one incipient attempt to probe Westin's privacy classification scheme and offer a critical analysis through empirical evidence of the predictability and generalizability of his work. Arguably, privacy researchers should be realistic about the limits of Westin's scale and be judicious in how they apply it. It might be that as we identify the political, socio-cultural, and personal factors that underlie privacy attitudes in the early twenty-first century, we find that Westin's fundamentals are reliable. But we also may discover other significant factors that contribute to privacy attitude formation.

I suggest it is time to stop using Westin's scale as the primary measure for understanding consumer privacy, especially outside of its original context, and pointedly, for explaining online information privacy attitudes. As my findings suggest, Westin's scale may give us little predictive capacity for assessing privacy knowledge, behaviors, or an alternative measure of attitudes. As we seek to understand the public's attitudes towards issues such as online cookie tracking, disclosure on social media, and other evolving privacy issues, Westin's scale likely can't provide us with a framework that gives us reliable predictive power for measurement. Rather than continuing to overload his scale, the privacy research community would do better to collectively develop a new instrument for reliable measurement.

7. REFERENCES

[1] Westin, Alan. Opinion Surveys: What Consumers Have To Say About Information Privacy. Testimony to the Subcommittee on Commerce, Trade, and Consumer Protection, May 8, 2001. Available at: http://archives.energycommerce.house.gov/reparchives/107/hearings/05082001Hearing209/Westin309.htm.

[2] Margulis, Stephen. On the Status and Contribution of Westin and Altman's Theories of Privacy. *Journal of Social Issues*, Vol. 59, No. 2, 2003, pp. 411-429.

[3] Turow, Joseph; King, Jennifer; Hoofnagle, Chris; Bleakley, Amy; and Hennessey, Michael. Americans Reject Tailored Advertising and the Three Activities That Enable It. September 29, 2009. Available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1478214

[4] Hoofnagle, Chris; King, Jennifer; Li, Su; and Turow, Joseph. How Different are Young Adults from Older Adults When it Comes to Information Privacy Attitudes and Policies? April 14, 2010. Available at: http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1589864.

[5] Westin, Alan F. Social and Political Dimensions of Privacy. *Journal of Social Issues*, Vol. 59, No. 2, 2003, pp. 431-453.

[6] Trepte, S., Dienlin, T., Reinecke, L. (in press). Risky behaviors—How online experiences influence privacy behaviors. In B. Stark, O. Quiring, & N. Jackob (Eds.), Von der Gutenberg-Galaxis zur Google-Galaxis. UVK: Wiesbaden.

[7] Cheshire, C., Antin, J. and Churchill, E. (2010), Behaviors, adverse events, and dispositions: An empirical study of online discretion and information control. J. Am. Soc. Inf. Sci., 61: 1487–1501. doi: 10.1002/asi.21346

[8] Jennifer King, Airi Lampinen, and Alex Smolen. Privacy: Is There An App For That? Presented at the Symposium on Usable Privacy and Security, July 2011. Pittsburgh, PA.

[9] National Criminal Justice Reference Service. Identity Theft – Facts and Figures. Visited May 15, 2014. https://www.ncjrs.gov/spotlight/identity_theft/facts.html

Appendix

Table 3: Demographic Variables

Variable	Type	Interpretation/Reference
Female	Binary	Male = 0 as reference cat.
Age	Continuous	18 is lowest value
Political views	Ordinal	Conservative (1) to very liberal (5)
Attends religious services	Ordinal	Attends more than once per week (1) to never (6)
Education	Ordinal	Less than high school education (1) to college degree and higher (4)
Race	Binary	White = 0 as reference cat
Household Income	Binary	Income below 50K = 0 as reference

Table 4: Privacy Knowledge Variables

Variable	**Type**	**Interpretation/Reference**
Online Privacy Questions	Ordinal	Summary of correct answers to Online Privacy Questions (0-5)
Offline Privacy Questions	Ordinal	Summary of correct answers to Offline Privacy Questions (0-4)
Self-reported internet skill measure	Ordinal	Beginner to advanced (1-3)

Table 5: Privacy Behavior Variables

Variable	**Type**	**Interpretation/Reference**
Erasing cookies	Binary	No = 0 as reference cat.
Read Privacy Policies	Binary	No = 0 as reference cat.
Respondent has refused to provide info to a business or website	Binary	No = 0 as reference cat.
Respondent is a member of a social networking website	Binary	No = 0 as reference cat.

Table 6: Privacy Attitude Variables

Variable	**Type**	**Interpretation/Reference**
Right to know everything a website knows about you	Binary	No = 0 as reference cat.
Right to ask for deletion of all personal data from a website or advertising company	Binary	No = 0 as reference cat.
Attitudes towards online targeted advertising (composite)	Ordinal	Average of responses to 5 online advertising questions, from less to more privacy protective (1-4)
Attitudes towards internet service providers sharing customer data	Ordinal	Less to more privacy protective (1-4)
Attitudes towards online advertiser data storage	Ordinal	Less to more privacy protective (1-4)

Table 7: Demographic Binomial Logistic Regression Results

Note: all results reported as odds ratios with linearized standard errors

	Female	Age	Pol. Views	Religious Services	Education	Race	Income
Fundamentalists	.87(.178)	1.02(.006)*	.92(.109)	1.09(.068)	1.07(.121)	1.40(.341)	.950(.209)
Pragmatists	1.02(.183)	.993(.005)	1.19(.111)	.997(.053)	.905(.088)	.802(.165)	1.34(.234)
Unconcerned	1.01(.227)	.969(.007)*	.884(.097)	.918(.064)	1.23(.171)	1.18(.318)	.666(.148)

*p=.000; Fundamentalists (F=2.21, p=.032); Unconcerned (F=4.13, p=.000)

Table 8: Privacy Knowledge Binomial Logistic Regression Results

Note: all results reported as odds ratios with linearized standard errors

	Online Privacy Questions	Offline Privacy Questions	Internet Skill
Fundamentalists	1.14(.076)*	1.20(.082)	.960(.141)
Pragmatists	.984(.058)	.971(.054)	.956(.110)
Unconcerned	.893(.084)	.906(.073)	1.33(.221)

*p<.05; Fundamentalists (F=5.56, p=.001)

Table 9: Privacy Behavior Binomial Logistic Regression Results

Note: all results reported as odds ratios with linearized standard errors

	Deletes Cookies	Reads Privacy Policies	Refused to Provide Info	SNS Member
Fundamentalists	1.02(.249)	.821(.206)	1.85(.714)	.764(.147)
Pragmatists	1.05(.209)	.957(.192)	1.00(.257)	.890(.135)
Unconcerned	1.06(.324)	1.38(.416)	.648(.201)	2.10(.437)*

*p<.000; Unconcerned (F=3.67, p=.006)

Table 10: Privacy Attitudes Binomial Logistic Regression Results

Note: all results reported as odds ratios with linearized standard errors

	Right to Know	Right to Delete	Targeted Ads	ISP Share	Ad Storage
Fundamentalists	1.73(.421)*	.989(.450)	1.33(.295)	1.27(.269)	1.44(.262)**
Pragmatists	1.14(.201)	1.12(.365)	1.12(.183)	.936(.138)	.780(.087)***
Unconcerned	.527(.119)*	.779(.294)	.713(.145)	.870(.170)	.926(.120)

*p<.01; ** p=.05;***p<.05; Fundamentalists (F=3.73, p=.002);Unconcerned (F= 4.14, p=.001)

Table 11: Demographic Multinomial Logistic Regression Results

Note: The Privacy Unconcerned are the reference category

	β_0	Female	Age	Pol. Views	Religious Services	Education	Race	Income
Fundamentalists	-2.30(.754)*	-.145(.278)	.045(.009)*	.923(.150)	.134(.083)	-.103(.166)	.143(.324)	.267(.277)
Pragmatists	-.263(.662)	.026(234)	.023(.007)*	.221(.122)	.067(.074)	-.201(.145)	-.222(.280)	.436(.231)

*p<.005; F=2.78, p=.001

Table 12: Privacy Knowledge Multinomial Logistic Regression Results

Note: The Privacy Unconcerned is the reference category

	β_0	Online Privacy Questions	Offline Privacy Questions	Internet Skill
Fundamentalists	.031(.492)	.200(.107)	.223(.097)*	-.266(.200)
Pragmatists	1.46(.435)**	.089(.098)	.068(.083)	-.249(.168)

*p<.05; **p<.001; F=3.57, p=.002

Table 13: Privacy Behavior Multinomial Logistic Regression Results

Note: The Privacy Unconcerned is the reference category

	β_0	Deletes Cookies	Reads Privacy Policies	Refused to Provide Info	SNS Member
Fundamentalists	.143(.608)	-.042(.361)	-.425(.363)	.856(.449)	-.829(.608)*
Pragmatists	1.42(.490)*	-.036(.314)	-.290(.310)	.344(.323)	-.664(.215)*

*p<.005; F=2.04, p=.040

Table 14: Privacy Attitude Multinomial Logistic Regression Results

Note: The Privacy Unconcerned is the reference category

	β_0	Right to Know	Right to Delete	Targeted Ads	ISP Share	Ad Storage
Fundamentalists	-5.01(1.25)	.965(.304)**	.198(.543)	.534(.272)***	.303(.267)	.379(.210)
Pragmatists	-.727(.695)	.581(.236)*	.247(.397)	.326(.209)	.076(.196)	-.038(.132)

*p=.002; ** p<.001; ***p=.05; F= 3.07, p=.001

Table 15: Mean age for each Westin category

	Age (mean)	Linearized standard error
Fundamentalists	49	1.5
Pragmatists	44	.81
Unconcerned	36	1.3