

A “Nutrition Label” for Privacy

Patrick Gage Kelley

Through an iterative design process, including focus groups and a laboratory study, we developed a standardized, tabular, “nutrition label” for online privacy policies. We tested the standardized format, two variants, and two real-world policy formats in a large, online user study to show that this label helps consumers.

Introduction & Motivation

Website privacy policies are intended to *assist* consumers. By notifying them of what information will be collected, how it will be used, and with whom it will be shared, consumers are, in theory, able to make informed decisions. These policies are also meant to inform consumers of the choices they have in managing their information: whether use of their information or sharing with third parties can be limited, and if it is possible to request modification or removal of their information.

However, policies are commonly long, textual explanations of data practices, most frequently written by lawyers to protect companies against legal action. It has been established through numerous studies that people do not read privacy policies [24] and make mistaken assumptions based upon seeing that a site has a link to a privacy policy [30]. A recent study estimated that if consumers were somehow convinced to read the policies of all the companies they interact with, it would cost an estimated 365 billion dollars per year in lost productivity [23]. In addition, research has shown that consumers do not actually believe they have choices when it comes to their privacy. Based solely on expectations, they believe there are no options for limiting or controlling companies’ use of their personal information [19].

In short, today’s online privacy policies are failing consumers because finding information in them is difficult, consumers do not understand that there are differences between privacy policies, and policies take too long to read. We set out to use a machine-readable standard and a firm HCI methodology to design a clear, uniform, single-page summary of a company’s privacy policy that would help to remedy each of these concerns.

Related Work

To better inform our design process we surveyed the literature surrounding other consumer labeling efforts: the Nutrition Facts panel, energy and drug labeling, and recent work on designing a financial privacy notice. Additionally, we summarize our previous work on a standardized privacy policy format.

Standardized Labels

In the United States, the nutrition label seen at right, has become iconic after being mandated by the Nutrition Labeling and Education Act of 1990 (NLEA) [32,33]. The sparse

literature around the design of the nutrition label [3] focuses on the decisions made to simplify the information as much as possible for consumers, in part to address low literacy rates and the needs of older Americans. These guidelines include defining a zone of authority, providing quantitative information about nutrients, defining minimum font sizes, and equalizing labels across products through serving sizes and calculating percentages based on standardized daily amounts.

Studies conducted to examine the impact of the NLEA have found that it is the people who are educated and already motivated to investigate nutritional information who benefit the most from nutrition labels [2][10]. Another study found that nutrition information had the greatest impact when there was a limited number of items from which to make a selection [27].

We also explored energy labeling programs from the European Union [13] and Australia [12], the US Consumer Products Safety Commission's toy and game warnings [8], and the US FDA Drug Facts label [32], to gain a broader understanding of practices used in designing and defining labeling requirements. In general, the standards documents [8,12,13,14,34] are occupied with defining precise guidelines to describe compliance with the various labeling requirements. This includes point sizes of rules and text, allowable typefaces, allowable colors, and minimum sizes.

**Sample Label for
Macaroni and Cheese**

Nutrition Facts	
Serving Size 1 cup (228g)	
Servings Per Container 2	
Amount Per Serving	
Calories 250	Calories from Fat 110
% Daily Value*	
Total Fat 12g	18%
Saturated Fat 3g	15%
Trans Fat 1.5g	
Cholesterol 30mg	10%
Sodium 470mg	20%
Total Carbohydrate 31g	10%
Dietary Fiber 0g	0%
Sugars 5g	
Protein 5g	
Vitamin A	4%
Vitamin C	2%
Calcium	20%
Iron	4%

* Percent Daily Values are based on a 2,000 calorie diet. Your Daily Values may be higher or lower depending on your calorie needs:

	Calories:	2,000	2,500
Total Fat	Less than	65g	80g
Sat Fat	Less than	20g	25g
Cholesterol	Less than	300mg	300mg
Sodium	Less than	2,400mg	2,400mg
Total Carbohydrate		300g	375g
Dietary Fiber		25g	30g

Nutrition Label from the Food and Drug Administration fda.gov

Financial Policy Notices

The Gramm-Leach-Bliley Act (GLBA), passed in 1999, contains the Financial Privacy Rule, which requires that financial institutions disclose their privacy policy “at the time of establishing a consumer relationship...and not less than annually” [31]. Financial institutions must comply with requirements on what they disclose, but their disclosures may be in any format.

In 2004, seven federal agencies launched a multi-phase initiative to “explore the development of paper-based, alternative financial privacy notices...that are easier for consumers to understand and use” [19]. The Kleimann Communication Group (KCG) conducted the first-phase, which tested multiple designs across seven cities and surveyed consumers about financial privacy notices. In their report the KCG proposed a three-page design for further evaluation [19]. In December 2008, the second phase report was published by Levy and Hastak [20]. This report detailed a 1032-participant mail/interview study that tested four privacy notice formats. Levy and Hastak concluded that the KCG table notice performed best. They attributed this improvement to an increased level of comprehension, given the table notice’s “[provision] of a fuller context...the part-to-whole display approach seems to help consumers focus on information sharing as important and differentiating features of financial institutions.” However, on several study questions other notices, notably the sample clause notice, tested best.

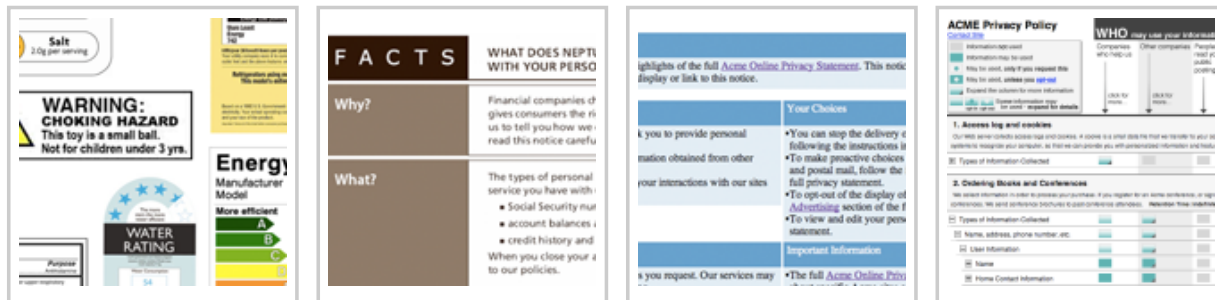
Layered Policy Notices

Layered privacy notices, popularized by the law firm Hunton & Williams [28,29], provide users with a high-level summary of a privacy policy. The design is intended to be a “standardized” format; however, the only standard components are a tabular page layout and mandatory section headers. Other design details and the text of each section are left to the discretion of each company. Additionally, the amount of information to include is left up to each company, requiring consumers to click through to the full policy.

The Platform for Privacy Preferences

Due to the difficulties surrounding the use of text privacy policies, the World Wide Web Consortium created the Platform for Privacy Preferences (P3P) [35]. P3P is a standard machine-readable format for encoding the online privacy policy of a company or organization. Once a P3P policy has been provided, consumers must use a user-agent to interpret it into something understandable. Unfortunately, widely available P3P user agents have limited functionality. These include the P3P policy processing elements of common web browsers and a few privacy specific browser add-ons [7].

To provide consumers with an active tool where they can investigate and explore the privacy policy of a website, earlier work from the CyLab Usable Privacy and Security Lab (CUPS) produced the P3P Expandable Grid. This user agent was based on one of the central Expandable Grid objectives of displaying a holistic policy view [25]. The interface was created to use the entire P3P specification, broken down by categories. However, based on an online survey of over 800 people in the summer of 2007, we found that people generally do not understand the information presented in privacy policies and also do not enjoy reading them. When comparing a natural language policy and the P3P Expandable Grid, we found neither of the formats were pleasurable to read or easy to comprehend. Notably, we found the P3P Expandable Grid to be worse than text policies, both in enjoyment and comprehension [26].



To view examples of these other formats and labels click the thumbnails above.

A Unique Approach

Our first goal was to design a "Privacy Label" that is actually understandable, allows users to quickly and efficiently find information, makes comparisons easy, and makes the experience of reading a privacy policy more enjoyable. Our second goal was to test in a large scale online study that our design met these goals.

Iterative Design Approach


Here we describe our iterative design process. While many designs were tested and eventually re-factored or abandoned, each of the examples given below show one of many variants of a similar vein. We have selected examples that we believe are representative of the major stages and changes throughout the process.

Problems with the P3P Expandable Grid and Labeling Insights

Through reviewing the P3P Expandable Grid study results and hosting a subsequent lab evaluation we identified five major problems with the Expandable Grid [18], most of which concerned providing users with far too much information in a hierarchy they were not familiar with. With these problems in mind we abstracted several general principles from the nutrition labeling literature [1,3,11,32,33]. For example, putting a box around the label identifies the boundaries of the information, establishing a zone of trust. Other common design elements involved using bold rules to separate sets of related information and

providing a clear and boldface title to communicate the label's purpose. While much of the labeling literature focuses on quantifiable properties, such as amounts of fats or fiber, or percentages of active ingredients or calories from a standardized expected daily value, privacy policies typically do not include quantifiable measures, and the P₃P specification includes no quantifiable fields. For financial privacy, the Kleimann Group dealt with this lack of quantifiable information by moving to binary Yes/No statements, which they found to be readily understood by focus group participants.

The Design Evolution



The Simplified Label

The Simplified Grid

The Privacy Label

Symbols

Visual Intensity

Useful Terms

The illustrations above give a short overview of the features and evolution of our design, here we describe several of the design contributions.

Simplification While we made visual changes including adding a title and subhead, adding bold lines, and simplifying the table view to create our simplified label, the most significant change is a reduction in complexity from the P₃P Expandable grid. Two changes contributed most to simplifying the label: eliminating P₃P statement groupings and eliminating the use of P₃P data hierarchies. The true depth of these changes requires a significant understanding of the structure of P₃P, so for a complete explanation please see [17].

Reintroducing The Table While the P₃P Expandable Grid was not successful, this failure was likely not a result of the tabular display. Also, due to the nature of P₃P Statements, each reduction in dimensionality causes a loss of information. With the reintroduction of the two

dimensional layout several other changes were made, including the further grouping of purposes and data types. We grouped the 12 purpose elements into six groupings, and merged the 17 data categories into nine groupings, as shown in the example. For information sharing, we show only sharing with other companies and sharing in public fora.

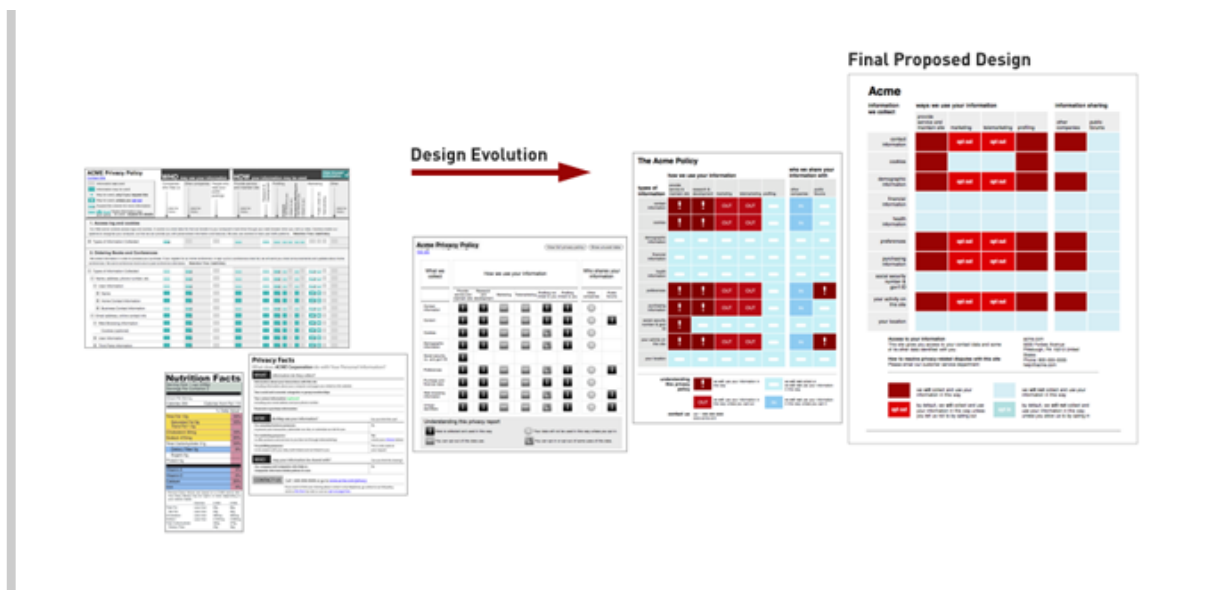
Symbols While you cannot opt-in or out to the trans-fat in your salad dressing, you might have control over certain aspects of your information sharing on the internet. The Yes/No dichotomy advocated by participants in the Kleimann Group's studies works when there are only one, or maybe two, columns of information. Here we would have needed 8 columns and 10 rows of Yes/No information, which would have been visually difficult to parse. Instead we again looked back to the P3P Expandable Grid and used the symbols shown above. Where the P3P Expandable Grid had an array of 10 symbols, the simplified grid uses only four, (and a blank symbol). Each of these four symbols is defined in a legend labeled "Understanding this privacy report" directly below the policy.

Visual Intensity The simplified grid is also the first iteration of our label to use visual intensity to provide a high-level indication of the quality of a given policy. Each of the four symbols has been colored such that darker symbols represent what could be more privacy-invasive practices. The use of intensity allows users to make quick visual comparisons that are not possible with text alone.

Blank Spaces In the simplified grid design, we marked types of information that companies collected and left other cells in the policy blank. However, half of the participants were actually worried by the blank spaces; for instance, one said, "Nothing is mentioned. It is completely open-ended. These guys [the company] can modify these values." Therefore, in the final version we introduced a symbol to indicate that information was not collected or used. Focus group participants found the mixed choices symbol confusing so we removed it. Instead we now display the symbol for the most invasive practice. For example, if in some circumstance one can opt-in and in another one can opt-out, we display the opt-out symbol.

Adding Color We originally constrained our initial designs to grayscale to facilitate easy printing without loss of information and to reserve color for highlighting differences between a policy and a user's personal preferences (something we plan to implement later). However, feedback indicated that color improved user self-reported enjoyment in reading the label. We selected the colors used in our label with care to accommodate viewers with color-blindness, allow for grayscale reproduction, and maintain the darker-is-worse high-level visual feedback.

Useful Terms Even with the legend in place there was still confusion over many of the terms used in the label. This was also a common issue during the development of the Kleimann Group's Financial Privacy Notice, and in response they developed what they call the "Secondary Frame." This portion of the prototype notice included both frequently asked questions and a series of extended definitions, which are: "[not] information as essential for consumers to have, but consumers often commented that they liked having it included." [19 p.27] Our useful terms information was informed by the Human Readable definitions included in the P3P 1.1 Working Group Note [36] and consists of seventeen definitions, one for each of the row and column headers.



Focus Groups & Qualitative Feedback

We held four, hour-long focus group sessions to review the design and discuss participants' impressions and questions. We recruited focus group participants from the Carnegie Mellon University (CMU) Center for Behavioral Decision Research (CBDR) participant recruitment website for the first three, and the final focus group was held at the Pittsburgh Jewish Community Center. We paid participants from CBDR \$10 to participate in our 60 minute focus group.

Our focus groups compared all three designs above (the simplified grid, the simplified label, and the first privacy label) as well as full text policies. The participants reacted positively to the tabular formats. For example, one participant stated, "This is more convenient than scrolling through reams and reams of paragraphs. I mean who reads them?" and another participant said, "I like the chart. [It's] better than long sentences." However, we found that some participants still had problems understanding privacy concepts. For example, one participant asked, "What is the difference between opt-in or opt-out?" Additionally, while most participants were familiar with profiling, but did not understand the difference between "Profiling linked to you" and "Profiling not linked to you." It was this vein of feedback that led to the inclusion of the useful terms definitions.

When reviewing the privacy nutrition label vs. the simplified label we found that participants better understood the table and were able to make more accurate side-by-side comparisons. Participants understood the significance of the red symbols, saying, "Red is for 'stop' or 'danger.'" We passed out two privacy policies, Policy A and Policy B, and asked the participants to raise their hands if they believed that Policy A is the better policy. Every participant raised his or her hand, correctly identifying Policy A as the more favorable policy. Some participants in this group even noted subtle differences between the two policies saying, "Policy A isn't perfect either, because they share your preferences, and this may include things like your religious or political preferences."

Regarding the simplified label, participants felt that it did not provide enough information, saying, "This is an empty policy, it says nothing. I wouldn't trust it." Participants wanted to see how each piece of information was being used. For example, one participant stated, "With the grid it's easier to see things. What information is being shared? We don't know that anymore."

Participants at the Jewish Community Center represented a very different demographic from the mainly college age participants in the earlier discussions. One important insight from this group was that the terms opt-in and opt-out which many students were unfamiliar with were well understood by an older population due to their use in medical paperwork. This

finding led us to reinstate the labels opt-in and opt-out, after having removed the word “opt” based on student feedback.

User Laboratory Study

After the first two focus groups we performed a 24-participant laboratory user study comparing a standard natural language (NL) privacy policy with privacy policies presented in our privacy nutrition label.

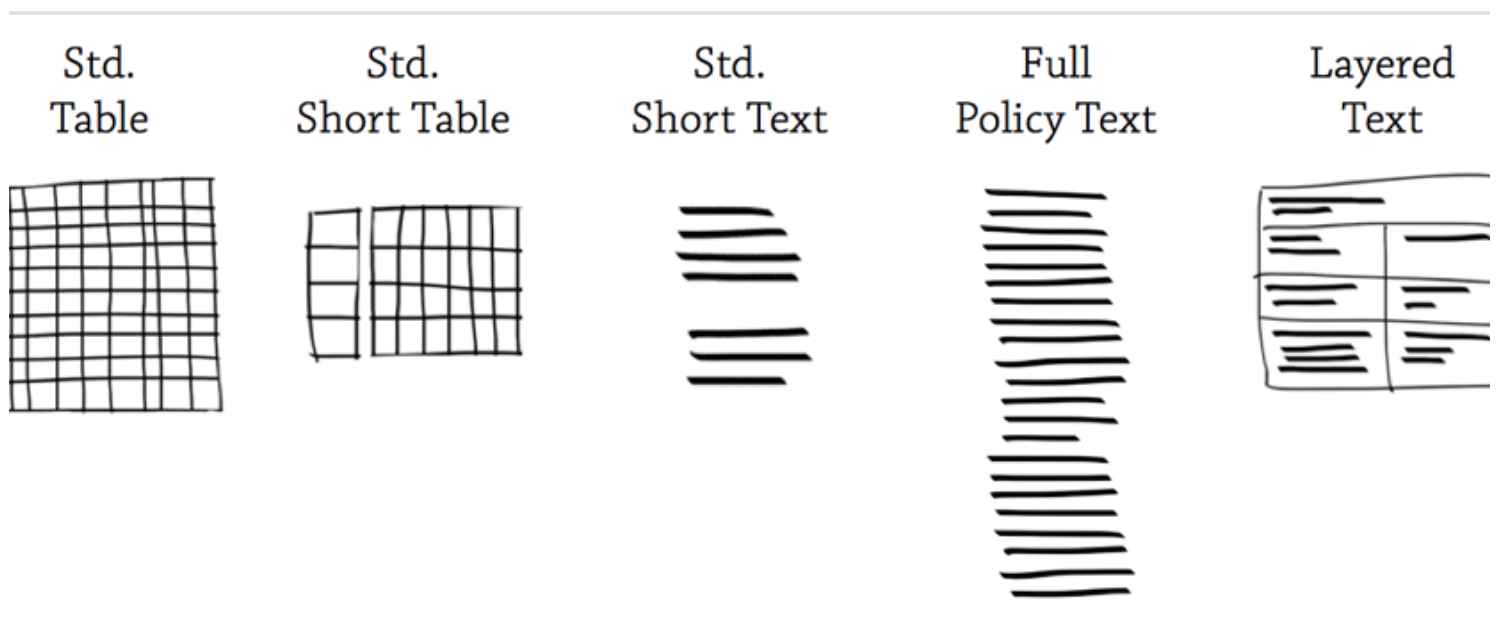
At a high level, people were able to answer more questions correctly with the label. We compared the correct number of total questions, per participant, for the label vs. the natural language policy, $M=10.13$ and $M=6.83$ respectively ($t(23)=7.41$, $p<0.001$).

The label was significantly faster than the natural language policies for both the group of information-finding questions and the group of policy-comparison questions ($p<0.001$). To test the mean task completion time for accuracy answers we removed all timing results where the answer was inaccurate and calculated means per question, per condition. Using a 2-sided t-test the label performed significantly faster in 2 of the 8 information-finding questions and 3 of the 4 policy-comparison questions.

All but 2 of the 10 Likert questions resulted in significant results. The label was rated significantly more pleasurable, easier to find information in, and easier and more enjoyable to use when comparing two policies.

Large-Scale Test

To follow up our design process and testing with a full-scale experiment, we tested five privacy policy formats: standardized table, standardized short table, standardized short text, full policy text, and layered text. Three of these formats are standardized and were created by our lab using an iterative design approach. Of these, two are tabular and one is textual. Two explicitly describe absent information and one presents it in the context of the policy. Each of these formats is followed by a list of definitions of privacy terms, consistent across formats. The policies are compared in the following graphic.



Pictographic representation of the final five formats compared.

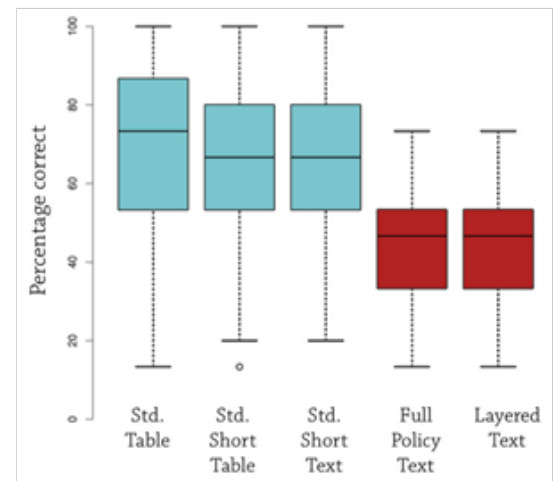
We conducted an online user study using Amazon’s Mechanical Turk, which offers workers the ability to perform short tasks and get compensated. For our approximately 15-minute study, we paid \$0.75 on successful completion. We developed a custom survey-management tool called Surveyor’s Point to facilitate our data collection. Our implementation allows us to show respondents a single question on the screen along with links for switching back and forth between two policies within a single browser window.

In preparation for this study we first performed three smaller pilot tests of our survey framework. We ran our pilot studies with approximately thirty users each, across 2-3 conditions. Our pilot studies helped us to finalize remaining design decisions surrounding the standardized short table, refine our questionnaire, and test the integration of Surveyor’s Point with Mechanical Turk.

We then conducted our large-scale study and completed the analysis with 764 participants (409 female, 355 male), randomly assigned to the five formats addressed above. We dropped 25 additional participants from the study prior to analysis due to incomplete data or for completing the study in an amount of time that indicated inadequate attention to the task (defined as time on task that was two standard deviations lower than the mean). We chose a between-subjects design to remove learning effects and ensure the study could be completed within about 15 minutes. Participants in each condition followed the same protocol; only the policy format differed.

Overall Accuracy Results

We scored each participant on a scale from 0 to 15, based on the number of the 15 information finding questions they answered correctly, and averaged those scores across conditions. Correct answers varied between conditions since policy content varied between conditions. We present these aggregate results on the right. This summary shows a large divide between the standardized and non-standardized formats (ANOVA significant at $p < 0.05$, $F(4,1094) = 73.75$). The three standardized formats, scoring 62-69%, are shown in light blue; while the two real-world text policies, scoring 43-46%, are shown in red. The standardized policies significantly outperformed the full-text policy (standard table vs. full text, $t(510) = -14.4$, standardized short table vs. full text $t(490) = 12.9$, and standardized short text vs. full text $t(491) = -14.3$, were all significant at $p < 0.05$). The layered format did not perform significantly differently from the full text policy ($p = 0.83$, $t(314) = -0.21$).



Overall accuracy results averaged for each of the five policy formats.

Note: For more information on accuracy results, including per question/condition results, and the full questions, please see the [appendix](#).

Timing Results

We examined completion times for the simple, complex, and comparison tasks, as presented in the table below. Time for comparison tasks includes both information-finding tasks and preference questions. We tested statistical significance using ANOVA on the log-normalized time information across policy formats. For each of these three groups of questions, as well as the overall study completion time there were statistically significant differences across policy formats ($p < 0.0001$ for questions 1-6, 7-12, 13-17, and overall). The standardized formats significantly outperformed the full policy text in overall time (standard table vs. full text, $t(348) = 5.36$, standardized short table vs. full text $t(327) = -6.01$, and standardized short text vs.

full text $t(329)=-4.55$, were all significant at $p<0.05$). The layered format was also significantly faster than the full text policy ($p=0.025$, $t(238)=2.25$). The standardized formats, on average were between 26-32% faster than the full text policy, and 22% faster than the layered text policy.

Average Timing Information (in seconds)

Question #	Std. Table		Std. Short Table		Std. Short Text		Full Policy Text		Layered Text		F-Statistic (dof)
	avg.	σ	avg.	σ	avg.	σ	avg.	σ	avg.	σ	
1-6	236	205	210	103	237	174	367	248	317	406	15.994 (4,756)
7-12	176	194	135	73	163	122	249	358	186	210	8.751 (4,756)
13-17	158	125	148	97	169	122	236	227	187	157	5.094 (4,756)
Full Study	912	572	852	407	938	515	1267	810	1089	768	11.273 (4,756)

Overall timing results, in seconds, averaged for each of the five policy formats, with standard deviations. Results are shown based on question group.

Enjoyability Results

For the most qualitative of our measures, we asked the participants how they felt about looking at privacy policies. For the single-policy tasks, participants across the board reported that they felt "confident in my understanding of what I read of Acme's privacy policy." The three comparison Likert questions show a larger preference for the standardized formats over the full policy text. The questions asked whether comparing two policies was "an enjoyable experience," was "easy to do," and if participants "would be more likely to compare privacy policies" if they were presented in the format they saw. The gap between the full policy text and the standardized formats widens from about half a point when looking at a single policy to as much as one and a quarter points after making comparisons.

While the layered text notice performed quite similarly to the full policy text in accuracy measures, we see a very different result in participants' feelings about using layered notices. The Likert scores for layered policies were not significantly different than the standardized-table format (1-6: $t(756)=1.57$, $p=0.115$; and 7-9 $t(756)=-1.48$, $p=0.138$).

The comments provided by participants at the end of the study provide insights into their enjoyment. Participants who saw the full policy text described privacy policies as "torture to read and understand" and likened them to "Japanese Stereo Instructions." On the other hand, participants in the standardized-format conditions were more complimentary...

"This layout for privacy policies is MUCH more consumer friendly. I hope this becomes the industry standard."

Results & Contributions

We began this work with the goal of designing a privacy policy format that would actually benefit consumers. A format that helps users accurately and quickly find information, makes comparison's between policies easier, and provides a more enjoyable experience. Our design approach allowed us to explore other efforts in standardizations, labeling, and designing privacy policy formats, while quickly and iteratively building a library of testable formats.

The final label design allows for information to be found in the same place every time. It removes wiggle room and complicated terminology by using four standard symbols that can be compared easily. It allows for quick high-level visual feedback by looking at the overall intensity of the page, can be printed, can fit in a browser window, and has a glossary of useful terms attached. People who have used it to find privacy information rated it as pleasurable. They not only rated it better than the natural language, but actually rated it enjoyable to use.

The three standardized formats that were designed with usability in mind performed significantly better across a variety of measures than the full-text and layered-text policies that currently exist online today. The large amount of text in full-text policies and the necessity to drill down through a layered policy to the full policy to understand specific practices lengthens the amount of time and effort required to understand a policy. Additionally, more complex questions about data practices frequently require reading multiple sections of these text policies and understanding the way different clauses interact, which is not an easy task. We have shown here that it is not solely table-based formats, but holistic standardization that leads to success. Our standardized policies left no room for erroneous, wavering, or unclear text, serving as a concise textual alternative to tabular formats.

The standardized formats performed the best overall, across the variety of the metrics we looked at. The accuracy, comparison, and speed results eclipse the results of the text formats in use today. Yet, while the accuracy with our standardized formats is better than guessing, there is still room for further study and improvement. Complex information-finding tasks and policy comparison tasks proved difficult. Future work should continue to concentrate on not just how to present policy information, but also on how to facilitate comparisons. Levy and Hastak recommend continuing to provide better education and context to help consumers make better decisions [20]. While our attached list of definitions is a start, framing the policy with contextual information and presenting comparisons in more useful ways would be productive directions to take future research in usable privacy policies.

Acknowledgments

The design team was led by Patrick Gage Kelley and included Joanna Bresee, Aleecia McDonald, Robert Reeder, Sungjoon Steve Won, and Lorrie Cranor. Thanks to Lucian Cesca, Cristian Bravo-Lillo, Robert McGuire, Daniel Rhim, Norman Sadeh, Clare-Marie Karat, and Janice Tsai.

This work was supported in part by U.S. Army Research Office contract DAAD19-02-1-0389 (“Perpetually Available and Secure Information Systems”) to Carnegie Mellon University’s CyLab, by NSF Cyber Trust grant CNS-0627513, by Microsoft through the Carnegie Mellon Center for Computational Thinking, FCT through the CMU/Portugal Information and Communication Technologies Institute, and the IBM OCR project on Privacy and Security Policy Management.

References

1. Balasubramanian S. and C. Cole. Consumers' search and use of nutrition information: The challenge and promise of the nutrition labeling and education act. In *Journal of Marketing*, 2002.
2. Beard, T.C., Nowson, C.A., Riley, M.D. "Traffic-light food labels." *Med J Aust.* 2007;186:19.
3. Belser, B. Designing the Food Label: Nutrition Facts. *AIGA Journal.* 1994.
4. Buckley, P. and Shepherd, R. Ergonomic factors: The clarity of food labels. *British Food Journal.* 1993.
5. Byrd-Bredbenner, C., Alfieri, L., Wong, A., and Cottee, P. The Inherent Educational Qualities of Nutrition Labels. *Family and Consumer Sciences Research Journal*, Vol 29,

No 3, March 2001 265-280.

6. Cranor, L., Egelman, S., Sheng, S., McDonald, A., and Chowdhury, A. P3P Deployment on Websites. *Electronic Commerce Research and Applications*, Volume 7, Issue 3, Autumn 2008, Pages 274-293.
7. Cranor, L.F. *Web Privacy with P3P*. O'Reilly and Associates, Sebastopol, CA, 2002.
8. Consumer Product Safety Commission. "Labeling Requirements for Toy and Game Advertisements." 2008. <http://cpsc.gov/library/foia/foia08/brief/toygameads.pdf>.
9. DeJoy, D.M., Cameron, K.A., and Della, L.J. Post-exposure evaluation of warning effectiveness: A review of field studies and population-based research. *The Handbook of Warnings*. 2006. (35-48).
10. Downs J.S., Loewenstein G., and Wisdom J. Strategies for Promoting Healthier Food Choices. *American Economic Review*. 2009, vol. 99, issue 2, pages 159-64
11. Drichoutis AC, Lazaridis P, Nayga RM. 2006. Consumers' use of nutritional labels: a review of research studies and issues. *Acad Marketing Sci Rev*, no. 9.
12. The Energy Label. 2007. www.energyrating.gov.au.
13. European Union Commission Directive 98/11/EC "Energy Labeling." 1998. <http://eurlex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:1998:07:1:0001:0008:EN:PDF>.
14. Food Standards Agency. "Signpost Labeling Research." 2005 <http://www.food.gov.uk/foodlabelling/signposting/signpostlabelresearch/>.
15. Jensen, C. and Potts, C. Privacy policies as decision-making tools: an evaluation of online privacy notices. *SIGCHI*. 2004.
16. Kelley, P.G., L. Cesca, J. Bresee, and L. Cranor. Standardizing privacy notices: An online study of the nutrition label approach. Technical Report CMU-CyLab-09-014, Carnegie Mellon University, November 2009.
17. Kelley, P.G., J. Bresee, L. F. Cranor, and R. W. Reeder. A "Nutrition Label" for Privacy. In *Proceedings of the 2009 Symposium On Usable Privacy and Security (SOUPS)*, 2009.
18. Kelley, P.G., A. McDonald, R. Reeder, and L. Cranor. P3P Expandable Grids. Poster at Privacy MindSwap Carnegie Mellon University. 2007. <http://cups.cs.cmu.edu/soups/2008/posters/kelley.pdf>.
19. Kleimann Communication Group, Inc. Evolution of a Prototype Financial Privacy Notice. February 2006. <http://www.ftc.gov/privacy/privacyinitiatives/ftcfinalreport060228.pdf>.
20. Levy, A. and Hastak, M. Consumer Comprehension of Financial Privacy Notices. December 2008. <http://www.ftc.gov/privacy/privacyinitiatives/Levy-Hastak-Report.pdf>.
21. Maubach, N., Hoek J. "The Effect of Alternative Nutrition Information Formats on Consumers' Evaluations of a Children's Breakfast Cereal" *Proceedings of the EPartnerships, Proof and Practice - International Nonprofit and Social Marketing Conference* 2008.
22. McDonald, A., Reeder, R.W., Kelley, P.G., and Cranor, L.F. A Comparison of Online Privacy Policies and Formats. *Privacy Enhancing Technologies* 2009.
23. McDonald, A, and Cranor, L. The Cost of Reading Privacy Policies. *Telecommunications Policy Research Conference*, 2008.
24. Privacy Leadership Initiative. *Privacy Notices Research Final Results*, November 2001, <http://www.understandingprivacy.org/>.
25. Reeder, R.W. *Expandable Grids: A user interface visualization technique and a policy semantics to support fast, accurate security and privacy policy authoring*. PhD thesis, Carnegie Mellon. 2008. <http://www.robreeder.com/pubs/ReederThesis.pdf>
26. Reeder, R., Cranor, L., Kelley, P., and McDonald, A. A User Study of the Expandable Grid Applied to P3P Privacy Policy Visualization. *Workshop on Privacy in the Electronic Society*. 2008
27. Seymore, J.D., Lazarus Yaroch, A., Serdula M., Blanck, H.M., and Khan, L.K. "Impact of nutrition environmental interventions on point-of-purchase behavior in adults a review." *Preventative Medicine* 2004. 29: S108-S136.
28. The Center for Information Policy Leadership. *Multi-Layered Notices Explained*, 2004.

[http://www.hunton.com/files/tbl_s47Details/FileUpload265/1303/CIPLAPEC Notices White Paper.pdf](http://www.hunton.com/files/tbl_s47Details/FileUpload265/1303/CIPLAPEC_Notices_White_Paper.pdf).

29. The Center for Information Policy Leadership. Ten steps to develop a multilayered privacy notice, 2005. [http://www.hunton.com/files/tbl_s47Details/FileUpload265/1405/Ten Steps whitepaper.pdf](http://www.hunton.com/files/tbl_s47Details/FileUpload265/1405/Ten_Steps_whitepaper.pdf).
30. Turow, J. Feldman, L., and Meltzer, K. Open to Exploitation: American Shoppers Online and Offline. The Annenberg Public Policy Center. 2005. <http://www.annenbergpublicpolicycenter.org/NewsDetails.aspx?myId=31>.
31. United States Code. 6803. Disclosure of institution privacy policy, 2008. <http://www.ftc.gov/privacy/glbact/glbsub1.htm#6803>.
32. U.S. Food and Drug Administration. A Food Labeling Guide. Center for Food Safety & Applied Nutrition. 1999. <http://vm.cfsan.fda.gov/%7Edms/flg-toc.html>.
33. U.S. Food and Drug Administration. "Guide to Nutrition Labeling and Education Act Requirements" 1994. http://www.fda.gov/ora/inspect_ref/igs/nleatxt.html.
34. U.S. Food and Drug Administration. "New OTC Drug Facts Label" FDA Consumer Magazine. 2002. http://www.fda.gov/FDAC/features2002/402_otc.html.
35. W3C. The Platform for Privacy Preferences 1.0 (P3P1.0) Specification. <http://www.w3.org/TR/P3P/>.
36. W3C. The Platform for Privacy Preferences 1.1 (P3P1.1) Specification. <http://www.w3.org/TR/P3P1/>.
37. WELS Regulator (Australian Government). "WELS and Watermark." 2005. <http://www.waterrating.gov.au/compliance.html>.

© 2007-2010 Patrick Gage Kelley patrickgage.com & The CyLab Usable Privacy & Security Laboratory cups.cs.cmu.edu