User-Centered Privacy Communication Design

Margaret Hagan

Stanford Law School/d.school
559 Nathan Abbott Way
Stanford, CA 94305
mdhagan@stanford.edu

ABSTRACT

In this paper, we describe a user-centered privacy policy design project that we undertook at Stanford Legal Design Lab, in order to generate new models of business-to-consumer communications around data privacy. From our preliminary user research, rapid prototyping and testing, and refinement of new privacy communication designs, focused on a very particular archetype – a 20-40 year old who is tech-savvy yet largely disinterested in privacy policies -- we propose a series of new concept designs for technology companies to use when presenting the terms of their privacy policies, as well as general principles to ensure that communication of these policies are more engaging and actionable to this type of target user.

1. INTRODUCTION

The topic of privacy policy design has been subject to great analysis and design work over the past decades. But there is not yet a clear direction about how technology companies, and their teams of lawyers, privacy specialists, and security specialists, can best communicate privacy policy terms to their prospective and existing consumers. The standard privacy policy, composed as a legal policy document or consumer contract, in a series of paragraphs of legal text, still prevails.

At Stanford Legal Design Lab (a cross-disciplinary initiative between the Law School and the Institute of Design), we conducted a ten-week class project, "Legal Design Lab: Consumer Contracts", to focus on new modes of communicating privacy terms to laypeople. In particular, our focus was on two stakeholders: the lawyers at technology companies who aim to make their policy more engaging and comprehensible, and the layperson who uses mobile technology and does not currently read privacy policies. In particular, we chose to focus on end-users who are between 20-40 years old.

With a small cohort of twelve students, we used a user-centered design process to understand our stakeholders, choose specific types of laypeople to focus on, and then create and test new ways to engage them around privacy policy terms. Our goal was to surpass status quo design efforts, and identify promising new communication designs for our particluar user archetype. In this paper, we present the five key designs that emerged out of the process, that include a privacy dashboard, role model privacy preference implementation, story-based privacy disclosures, selective just-in-time alerts, and visual diagrams of where data goes and is stored. In addition, we recommend that future privacy policies prioritize more interactivity, modularity, and decentralizing the policy document into more discrete, relevant, and visualized communications that fit with how users experience the technology, rather than relegating all of the privacy terms to a single document that users will rarely engage.

Symposium on Usable Privacy and Security (SOUPS) 2016, June 22-24, 2016, Denver, Colorado.

2. BACKGROUND

This project emerged out of the proposition that design methods could improve the lay people's engagement with data privacy information, particularly concerning privacy of their data on their mobile phone. It developed in response to a growing debate in academic literature about whether companies' mandatory disclosures of terms and conditions could ever be effective, and whether better design could increase this effectiveness.

2.1 The potential power of design

In legal scholarship around consumer-facing contracts and government-mandated disclosures by companies to users, there is growing debate about whether design can effectively help users to engage with and understand the contents of the communications that companies publish for them.

Many contracts scholars propose that increased use of visual design principles and formats will make legal documents more accessible, comprehensible, and actionable. [7, 8]

Other scholars reject the notion that any amount of design could improve the effectiveness of 'mandatory disclosure' as a regulatory policy, because they hypothesize that laypeople will not engage with policy communications no matter how well composed or visualized. [2] Some of these scholars have attempted to impose better designs on privacy policies, then have studied laypeople's comprehension of these new designs, and found that the design changes that their team made did not have any meaningful effect. [1]

2.2 Previous Redesign Efforts

Our group was aware of the many privacy policy design projects that have come before, including speculative designs and implementations by companies and regulators.

These previous inventive designs include the expandable grid that visualizes privacy policy features and terms and the nutrition label for privacy policies. [5, 11] Another theme has been standardized

iconography, to signal either the types of privacy terms being discussed, or the actual practices described in these terms.[9, 10] These standardized icon sets correspond to the successful Creative Commons schema of labels and icons.[4] Another stream of design interventions feature more sensory aspects, like the Privacy Bird tool, which would have a browser-bar icon that changed colors and make a sound depending on how the policy of the site differed from the preferences the user expressed. [3]

We also were conscious of disclosure design efforts from analogous areas. Credit card terms and conditions' standardized disclosure design, created by several federal agencies, requires that all financial companies offering credit cards to consumers present the terms in a standard, table format, with larger text and key terms highlighted. [6] One company redesigned how healthcare providers doing medical research present the terms of participating in a study to laypeople, using iconography, limited text, and staging of information among a series of screens. [12]

2.3 Designs of current policies

Our group surveyed how technology companies are conveying their privacy terms and conditions, in order to observe what types of new designs and interactions are being used to improve engagement and comprehension of the disclosures.

Our class found distinct trends in how companies are trying to improve the presentation of their privacy policies.

Segmenting and Layering. Some policies segment the policy into many different titles and sections, in order to provide more navigation and of the content, and to direct users with specific interests to separate experiences while keeping the main policy relatively straightforward. (For example, Dropbox has sub-pages for specific topics, like government data disclosures).

Illustrated, step-by-step checkups. These communications exist on top of the formal policy, in order to present key terms or decision-points one-by-one, and to let the user take action to determine their preferences regarding some of the data privacy choices the site offers. These guides do not often present the company's policy in full, but rather makes the user aware of what choices the company does give them about their data privacy and then prompts the user to make more deliberate choices. (For example, Facebook's Privacy Dinosaur and Google's Privacy Health Checks)

Icons, white space, and navigation. Many policies improved the visual design of their policies by using larger fonts, providing short visual illustrations through the use of icons, allowing users to expand out content, and allowing users to jump to specific content through interactive navigation. (For example, the policies of Microsoft, Google, Yahoo, and Twitter).

Multi-user design. These policies are in some way modular, to adjust to different users. The policy is presented differently for regulators and legal experts, as compared to lay consumers. (For example, the policies of Codepen and Microsoft)

3. OUR USER-CENTERED DESIGN APPROACH

To create new concept designs, our group adopted a user-centered design approach. This involves narrowing the target audience for the new privacy design to specific types of individuals, rather than a generalized public, and creating and testing new design concepts in a rapid way. The expectation was that this approach would allow for greater creativity in the concepts we created, and greater responsiveness to the target audience's needs and preferences. This process, summarized below, led to the collection of new privacy business-to-consumer communication designs and guiding insights and principles for future privacy policy work.

3.1 Scoping the Challenge

We began our process with discussions with large technology company lawyers from a US company that provides both hardware and software mobile technology to lay people. With these lawyers, our team discussed how they currently craft privacy policies, and what their internal (inside their company) and external (public, media, etc.) stakeholders and concerns are. This helped us to understand what dynamics are at play for the lawyers, security experts, and communications professionals crafting the policies, and what their interests and goals are when deciding what other kinds of policies they would be willing to experiment with.

Our goal became to draft new concept designs that this and similar technology companies might implement voluntarily. We specifically intended to create designs that well-intentioned companies, which wish to give their users greater transparency about privacy practices and want more users to engage with their privacy policies, would customize and implement in place of, or in addition to, their current privacy communications.

We realize that this challenge scope does not cover all companies -- that some may prefer that their users not engage with or understand the privacy policy, and that not all are well-intentioned about the effectiveness of privacy communications. Our project did not focus on such companies, but those that do aspire for greater user engagement and comprehension. This challenge also was not focused on small start-up companies, though the designs and insights that emerged might be applicable to their privacy communications as well.

3.2 User Research and Defining Personas

After scoping down to the challenge of creating new privacy communications that mobile technology companies might implement, we then turned to understanding what possible users we might be targeting with these communications. We ran in-person interviews and online surveys in order to understand how our target users approached issues of privacy, and how they interact with privacy policies. Our in-person interviews had 20 participants, all of whom were young adults attending Stanford University as undergraduates or graduate students. Our online surveys had 100 participants from Mechanical Turk, all of whom own mobile phones, were between 20 and 40 years old, and ranked themselves as somewhat tech-savvy or above.

From this initial research, the group chose a particular target user upon which to focus their new privacy communication designs. For this user, the group created a user persona, which summarized the key preferences, information- and device-related behaviors, values, and aspirations about an archetypal version of their target user. This persona was not meant to capture wide demographics, but rather more particular types of people that the group had encountered during their research.

3.3 Brainstorm and Selection of New Concepts

Using these personas and design briefs, the group split into three teams of four, which brainstormed concept designs for new kinds of privacy communications for this target user on their mobile phone. The teaching team prepared the teams with examples of current types of privacy policies and experiences, so that they could understand what some of the current status quo trends are. The intention was to provoke the teams to go beyond these current models, and also abandon the constraints of the 'policy' on a 'privacy policy page'.

3.4 Testing and Refinement

The teams created a large collection of ideas, but then were asked to pare down to five concepts or less to prototype and test. They created rough mock-ups of each of these concepts, with sketches, storyboards, and short write-ups. They presented these early prototypes to their target users through in-person testing sessions on Stanford campus and online through surveys with 120 Mechanical Turk participants from the chosen demographic.

The prototypes were also subject to design reviews in our class, review from contracts scholars, and feedback from technology company lawyers and privacy experts. Their feedback complemented the users' expressed preferences, helping the teams understand if their concepts would be legally viable, and would be voluntarily implemented by the policy-makers inside the technology companies.

Based on the user, expert, and stakeholder feedback, the teams then refined their concept designs. They went from multiple ideas down to a handful, to create at a higher fidelity. For the chosen designs, they created full prototypes, with higher levels of visual design and a plan for how and where the design could be implemented. We present these designs here in the paper, as well as the design research and work we did to produce them, for further review and evaluation.

4. USER RESEARCH AND BRIEF

Our initial research into laypeople's relationship with privacy policies began with exploratory interviews, and then we moved from the insights we pulled out from the interviews to larger online surveys.

4.1 Exploratory Interviews and Online Surveys

Our exploratory interviews involved a small sample size of around 20 people, with whom we spoke for ten to twenty minutes in person. These in-depth discussions revealed several central insights into how young, technology-literate people interact with privacy policies. We then examined these insights with our online survey, which confirmed two key findings: users have a wide variety of privacy concerns that are not clearly prioritized or universal among all users in our chosen persona type; and that these concerns do not translate into engagement with a privacy policy document on their phone, inside apps, or on technology companies' websites.

Grab-bag of privacy concerns. First, there is a high variance about what people care about regarding privacy issues. Among the top eight themes, in no meaningful order, were: (a) avoiding spam content or messages; (b) preventing third party access to the data they share with a piece of software or device; (c) avoiding advertisements; (d) knowing if their private data is being monetized by others; (e) the security of their personal data, and how it was being stored; (f) knowing generally how companies are using their data; and (g) having a sense of control to respond and direct companies about their treatment of personal data.

Interest in privacy, but not privacy policies. The interviews also revealed that many people are thinking about issues of privacy, and weighing them as they consider how they interact with their mobile, device, but that they do not find the privacy policy to be useful enough to read. They engage in a 'cost-benefit analysis' when presented with the policy, and they ultimately decide that the time spent to read the policy and comprehend it would not be worth the value that reading the policy would give them. The people were not resigned to a loss of privacy -- in fact, they resisted the statement that "My information is out there, so privacy doesn't matter".

4.2 Target Users

The Privacy-Indifferent Young Professional: Our group chose to target a mid-20s young professional, who is comfortable with technology but does not use their mobile phone intensively. Primarily, they use the phone for consuming news and sharing items on social media. This persona has some exposure to what data privacy is, but does not have any specifically articulated concerns about their personal privacy vis-a-vis their mobile phone. Privacy is not a factor in how they purchase technology or use devices or software. They do not regularly (or ever) read privacy policies.

The team's design brief for this user was "How can we help this user, who is vaguely concerned about privacy but doesn't typically read privacy policies, to use and navigate information about their phone provider's privacy practices, so that they can understand the information that matters to them and take action (when possible)?"

5. CONCEPTS AND REFINED DESIGNS

After creating this design brief, the team then focused on creating a wide range of new privacy policy designs to test with our two key stakeholders – our target users and technology company policy-writers. The brainstorm produced a wide and ambitious selection of ideas, as well as some only incrementally different than some current companies' policy designs. These ideas (not including our final concept designs, listed later) included:

1. **Creating a Character for Personal Data**, that would help a user visualize and identify more with the data itself, and think more about it in concrete terms and less in abstract ones.
2. **Social sharing,** in which one's social network could flag parts of the privacy policy that concerned them or that they liked, so the user could better read the policy and know where to focus.
3. **Privacy reminders**, in which users would set timed alerts to remind themselves to review privacy terms and change the options they had selected.
4. **Conversation Bot,** that would let a user speak or type to a character (like Siri, Google Now, or other in-built bots from the phone provider) about privacy. The user could ask questions generally (e.g., "Hey Siri, Tell me about privacy") or about

specific topics (e.g., "OK Google, who do you share my location data with"), and the bot respond with details about the policy's terms, phrased in a conversational way.

5. **Icon-based Navigation**, that would have the policy grouped around certain topics, and have small, distinct illustrations to flag and detail these topics. It would allow users to easily scan the policy, jump to the topics they found relevant, and understand the core terms.

6. **Privacy Warnings and Mandatory Education**, that present roadblocks in the form of stop signs, mandatory education about the policy, and requirements to pass a short quiz before being allowed to accept terms or change a setting.

5.1 Feedback during user and expert testing

As we tested these ideas, many of them were left behind as too confusing, too demanding, or too emotional for laypeople. Company lawyers also resisted designs that were negative in tone, or that departed substantially from a privacy policy.

5.1.1 Focus on simplicity, low-requirement interventions

The brainstorm had produced many ideas for narrative- and game-based disclosures. These would ask users to explore the policy through a series of interactions, linked together with a narrative, characters, and scenarios. These designs did not test well, because they seemed inefficient and did not quickly deliver key information to the users. The testers preferred more of a disclosure that would give them core facts in ways that they could choose where to focus or not – but not be locked into a series of interactions that took substantial time and was a one-way track.

5.1.2 Preference for neutral, text communications

One concept design had three variations: one in which the text of the communication was presented in a "happy/upbeat" tone, a second in a "neutral" tone, and a third in a "fear and warning" tone. The team used both text phrases and visuals to convey these different tones. They presented all three variations to testers online and in-person.

The "neutral" tone received the most positive feedback. Participants preferred a text-centric design, but with the text in limited quantities and directly answering a question that is relevant to them. They preferred text-based communications, that included some visual iconography, over purely visual designs that the team had created, saying that the text provided more clarity when communicating complex terms about a topic they were otherwise unfamiliar with. They also recognized the emotional tones in the other design choices, and did not want the communication to be urging them to feel either upbeat or fearful. They wanted a communication that seemed to be free of value-judgment, so that they, the user, could assess the information and make a choice on their own.

5.2 Refined Design Proposals

The team concluded the project with a handful of concept designs that tested well with both laypeople and technology company lawyers. These five designs are the: (1) privacy choice dashboard; (2) the visual data privacy diagram; (3) the multi-character privacy stories; (4) the context-specific, simple alert; and (5) the role model preferences tool.

5.2.1 Privacy Choice Dashboard

One new proposal is for a central dashboard on a mobile phone, that would let a person view and control the privacy practices that apply to their data. For example, on the iPhone, there would be a distinct Privacy app, alongside the Settings app.

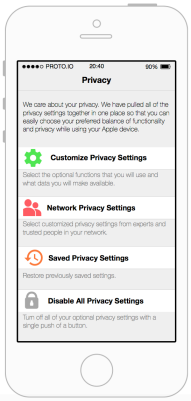

Figure 1. Privacy Dashboard main screen, for the phone's built-in 'Device and App Privacy' application.

On opening, the user could see what all of the different privacy trade-offs there are, based on the phone provider's policy, as well as third party app's policies. For any of these trade-offs that allow the user to make an opt-in or opt-out, the dashboard would let them make this choice immediately there. The dashboard would explain what the different groups of settings mean, and what the advantages and draw-backs of sharing data might be.

5.2.2 The Role Model Preferences Tool

The dashboard is connected to the second proposal: a way to instantly borrow and implement trusted others' privacy preferences.

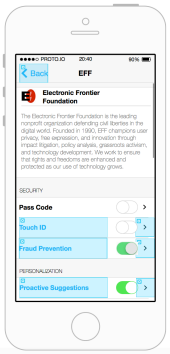

Figure 2. Using the EFF's privacy preferences as a role model.

This concept tool presents a range of possible role models – from well-known privacy groups, to prominent individuals, to people within one's social network – to implement their privacy choices on one's own dashboard. The user could save time and energy by adopting a curated set of data policy settings in one choice, rather than customizing the settings on their own. For example, the

privacy dashboard could offer them 'role models' or 'leaders' for privacy settings. A user could adopt these models' settings and tweak them.

5.2.3 Visual Data Privacy Diagram

A separate concept is a visualization of how data traveled, how it is shared, and how it is stored. They gathered together all the various policy clauses that applied to certain categories of data — like, for photo data, location data, etc. Then they created graphic maps of if this data stayed on the phones, on the mobile phone company's servers, on 3rd party servers, or even beyond.

One location for this visual diagram would be on the technology company's website, alongside or connected to their privacy policy documentation. A large desktop screen would be advantageous to see the entire diagram at once.

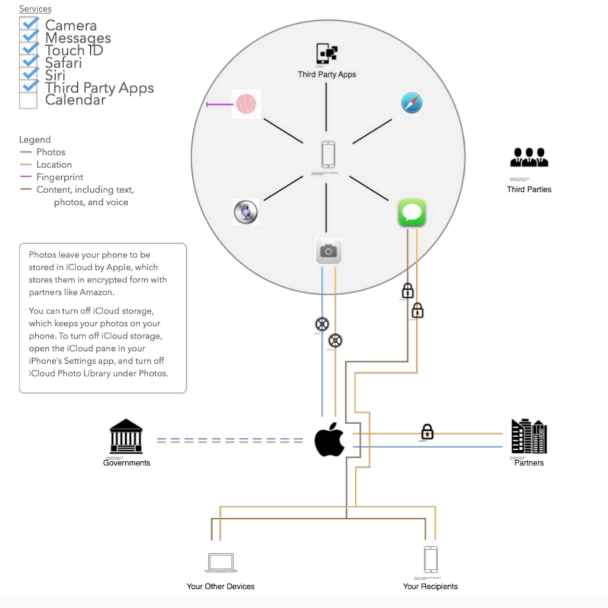

Figure 3. A Data Privacy diagram, visualizing where different kinds of user data is distributed and stored.

Another version of the diagram would be on the phone, on the app's setting or the policy page. Here the diagram would be radically simplified. It would only show the path of one kind of data, rather than the entire system of how different types of data (photos, location, etc.) travel.

5.2.4 Multi-character stories

A fourth concept is a story-based model. Rather than presenting information about the privacy practices through legal clauses and standard policy languages, the idea is to show the practices through a series of stories. The user could browse all the stories, or choose the one that seemed to correspond most to their own situation and demographics.

The design includes a series of fictionalized personas, based on the group's user research and testing. For each of these personas, we wrote out stories that demonstrated the core messages of the privacy policies, but through human examples. We curated the types of terms to mention, based on this type of user's concerns and preferences. The stories were told in a series of swipe-able screens, with small bits of text mixed with illustrations. The goal was to have a lightweight but engaging narrative to humanize the privacy terms, showing how they play out in concrete scenarios and what consequences they have for people.

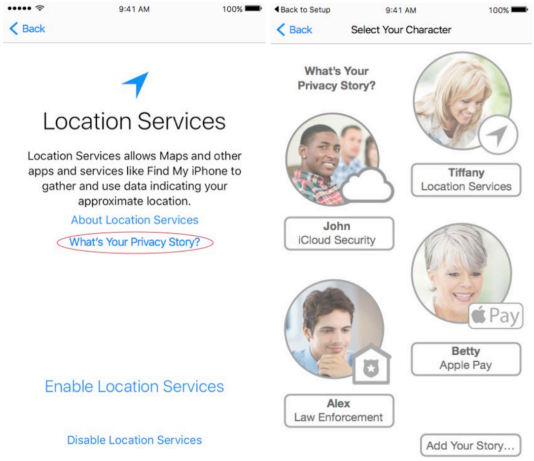

Figure 4. Privacy terms communicated through archetypal user stories, about privacy preferences, scenarios, and consequences.

5.2.5 The Context-Specific Alert

A final design proposal centered on one specific type of data privacy: that around location data. Phones or apps should have dashboards that pull out all the information that could affect how the user's personal data is treated. Particularly for users with strong interest in how their location data was used, this design focuses on selectively giving them these terms of the policies, mixed with actions they can take. These users didn't want to necessarily sift through the policy to find all the clauses that applied to location, and then figure out how it affected the treatment of their data. The goal was to present answers to their key questions about location data: where it was stored, and who has access to it, right when they are interested in these questions.

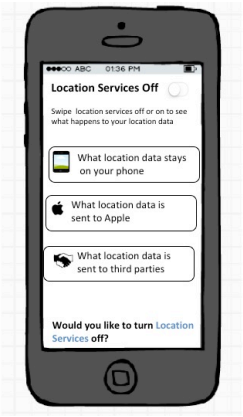

Figure 5. Context-specific location data terms alert.

This design has a streamlined location data privacy term explanation, along with the choice of turning this service on or off, at key 'privacy-aware' moments for the user. These moments include when the company has changed the terms, when a new

company begins to access location data, or when the user is electing to choose the location services on or off.

6. CONCLUSION

After reflecting on our design research and concept proposals, we identified a range of principles and a basic framework that that can guide future privacy communications. Along with the actual concept designs that we generated, these broader conclusions can be of use to scholars, technology company lawyers and government policy-makers, as they explore how to improve company-to-consumer privacy practice communication.

6.1 Principles for Improved Privacy Communication Design

6.1.1 Weaving the policy into the overall experience

A fundamental learning from our project was that the policy document should be broken up into smaller, more actionable communications that appear throughout the entire experience of the technology it applies to. The terms must be presented in more discrete and relevant ways. Future privacy policies can de-centralize the policy communication from a single page, covering all conditions and types of situations, and instead communicating specific types of terms right when that user might be interested or making a decision about that topic.

For example, the topics of third parties' use of location data that apps are gathering is often of concern. Policy terms about location may be shown both in the policy document, but also displayed right when a user is agreeing to share the location or when the location data is being collected. The terms here can be framed in terms of the shortlist of main questions a typical user might have while making this decision to share location data, or while observing the transfer of location data: What location data is sent to your mobile phone provider? What location data is sent to third parties? What location data is stored only on your phone and not shared? These contextual questions can help the user understand the policy terms in the frame of real choices and just-in-time.

6.1.2 Making human rather than legal disclosures

The communication of data privacy terms may serve a legal function, in creating a business-to-consumer contract and in showing regulators that the company has fulfilled its duty to its consumers to inform them of these terms. But the design of the communication need not take a traditional legal approach. It can be framed into narratives, with characters or a selection of personas. It can explain how the data privacy terms affect humans, what real-life scenarios may play out, and demonstrate how these terms may impact its users' choices, behavior, and privacy.

6.1.3 Illustrating trade-offs, showing what the user might 'lose'

Another theme that engaged our laypeople testers was presentation of the privacy policy terms in ways that show the consequences of the term, and the consequences if the user opts-out of the default or if the term doesn't exist. Essentially, this can show the user what services, features, and specific behaviors the user receives in exchange for giving their data to the company. Users report wanting a schematic that lets them evaluate terms so that they can easily weigh the advantages they will receive versus what they will give up, actually or potentially. They want to be smarter consumers, but they want an easy presentation of the trade-offs of their choices, so that they can quickly comprehend them and make the right decision for themselves.

6.2 Thinking Systematically about Privacy Communication Design

The design project led us to refine our conception of the types of communications that can be used to communicate privacy terms or other policy and legal disclosures. We present a framework of levels of disclosure design, from the most traditional to the most ambitious (and, ideally, the most effective). The aim here is to push beyond a cosmetic 'facelift' of a traditional policy document, and establish new types of communication.

A Level 1 type of disclosure design is a text-based document, that lists all the terms with minimal numberings, headings, and other structure to give the reader a basic sense of the different terms it contains. Some policy designs within this level can be better than others, using principles of plain language, that de-emphasize legal jargon, keep the length of sentences short, and exclude unnecessary or redundant information.

A Level 2 type of disclosure takes this simple text document, and provides more visual illustrations (through the use of icons, pictograms, tables, or charts) to offer the reader a more accessible presentation of the types of terms and their contents. It may also have a table of contents with links to the relevant portion of the policy, and other navigation features that allow the user to quickly jump through the text to the relevant portion. Essentially, the Level 2 of disclosure uses illustrations, navigation, and links to make the text document more accessible to the lay user.

Level 3 disclosures break up the text-document format of the communication, to instead present the terms with greater interactivity, staging, and visualization. This type abandons the lengthy policy document, and instead presents the user with a communication that she can selectively consume, and do so with more targeted, short presentations of information that the the user can interact with. These de-centralized disclosures are more just-in-time, more conversational, and more relevant to the user's concerns at that moment.

Another type of level 3 disclosure is not necessarily de-centralized, but reframed from a privacy policy to a privacy dashboard. If the choices and terms are taken out of the legal document, and instead presented in a Settings-style app, with choices integrated directly next to the terms of the policy, and all related policy terms linked together, then the user can have a centralized point of control over their data privacy on the device. This could even be a cross-app dashboard, that would allow a single point on the phone to control all first- and third-party apps' treatment of the user's data privacy. Our testing reveals that people are much more likely to engage with a dashboard of choices and policy terms than with a 'privacy policy' itself.

Encouraging technology companies to move towards this third level of policy design should be a priority going forward for scholars and policy-makers. This approach, of abandoning the formal legal policy document, can lead to new ways to make the terms meaningful and actionable for the lay users, and to overcome their resistance to engage with communications called "privacy policies."

7. REFERENCES

[1] Ben-Shahar, O. and Chilton A. (2016). "Simplification of Privacy Disclosures: An Experimental Test." *University of Chicago Coase-Sandor Institute for Law & Economics Research Paper No. 737*.

[2] Ben-Shahar, O. and Schneider, C.E., (2010). "The Failure of Mandated Disclosure," *U of Chicago Law & Economics, Olin Working Paper No. 516*.

[3] Cranor, L.F., Guduru, P. & Arjula, M. (2006). "User interfaces for privacy agents," In *ACM Trans. Comput.-Hum. Interact.* 13, 2, June, 135-178.

[4] Creative Commons, (2016). "CC-inspired projects for Terms of Service and Privacy policies," Available at https://wiki.creativecommons.org/wiki/CC-inspired_projects_for_Terms_of_Service_and_Privacy_policies, last accessed May 12, 2016.

[5] Kelley, P.G., Bresee, J., Cranor, L.F., & Reeder, R.W. (2009) "A 'nutrition label' for privacy," *In Proceedings of the Symposium on Usable Privacy and Security: SOUPS '09*.

[6] Kleinmann Group, (2009). "Web-based Financial Privacy Notice: Final Summary Findings Report," Available at https://www.ftc.gov/system/files/documents/reports/model-form-rule-research-report-creating-web-based-model-form/model_form_rule_research_report_on_creating_a_web-based_model_form.pdf , last accessed May 12, 2016.

[7] Mitchell, J. (2015) "Putting some product into work-product: corporate lawyers learning from designers," *In Berkeley Business Law Journal,* Vol. 12, no. 1.

[8] Passera, S., Haapio, H., & Barton, T.D. (2013). "Innovating Contract Practices: Merging Contract Design with Information Design," In California Western School of Law, CWSL Scholarly Commons, available at http://scholarlycommons.law.cwsl.edu/cgi/viewcontent.cgi?article=1072andcontext=fs, last accessed May 12, 2016.

[9] Raskin, A. (2009). "Making Privacy Policies Not Suck." Available at http://www.azarask.in/blog/post/making-privacy-policies-not-suck/, last accessed May 12, 2016.

[10] Raskin, A. (2010). "Is a Creative Commons for Privacy Possible." Available at http://www.azarask.in/blog/post/is-a-creative-commons-for-privacy-possible/, last accessed May 12, 2016.

[11] Reeder, R.W., Kelley, P.G., McDonald, A.M. and Cranor, L.F. (2008). "A user study of the expandable grid applied to P3P privacy policy visualization," *In Proceedings of the 7th ACM workshop on Privacy in the electronic society (WPES '08),* ACM, New York, NY, USA, 45-54.

[12] SAGE Bionetworks, (2016). "Participant Centered Consent Toolkit" Available at http://sagebase.org/platforms/governance/participant-centered-consent-toolkit/, last accessed May 12, 2016.