

CHICAGO

COASE-SANDOR INSTITUTE FOR LAW AND ECONOMICS WORKING PAPER NO. 776



COASE-SANDOR INSTITUTE
FOR LAW AND ECONOMICS

THE UNIVERSITY OF CHICAGO LAW SCHOOL

IS PRIVACY POLICY LANGUAGE IRRELEVANT TO CONSUMERS?

Lior Jacob Strahilevitz & Matthew B. Kugler

THE LAW SCHOOL
THE UNIVERSITY OF CHICAGO

September 2016

Is Privacy Policy Language Irrelevant to Consumers?

Lior Jacob Strahilevitz¹ & Matthew B. Kugler²

45 Journal of Legal Studies __ (forthcoming 2017)

Abstract

Consumers almost never read privacy policies, but if they did read such policies closely how would they interpret them? This article reports the results of two experiments in which census-weighted samples of more than a thousand Americans read short excerpts from Facebook, Yahoo, and Google's privacy policies concerning the use of facial recognition software and automated content analysis on emails. The question of what consumers have consented to under these policies has been central in recent high-stakes class action lawsuits. Experimental subjects were randomly assigned to read language from either the current policies, which explicitly describe Facebook, Yahoo, and Google's controversial practices, or language from policies that were adjudicated to be insufficient to notify consumers about the companies' practices. Despite evidence that many experimental subjects read these privacy policy excerpts closely, subjects who saw the explicit policy language and those who saw the ambiguous / vague policy language did not differ in their assessment of whether their assent to that language would allow Facebook, Yahoo, and Google to engage in the practices at issue. More surprisingly still, even though consumers rated both Facebook's use of facial recognition software and Google and Yahoo's use of automated content analysis as highly intrusive, they generally regarded their assent to even vague privacy policy language as allowing the companies to engage in those practices. Also, only a little more than a third of the participants expressed a willingness to pay any money to avoid automated content analysis of their emails. A replication study that included strong measures of participant attention confirmed the results from the first experiment and suggests that those reading the policies more carefully were not more likely to draw distinctions between them.

Our study shows that courts and laypeople can understand the same privacy policy language quite differently. Taken together, these results provide important evidence for the propositions that (1) social norms and user experiences with technological applications, not privacy policies, will drive users' understanding of the nature of their bargain with firms, that (2) this is the case even when users read those policies reasonably carefully, that (3) most users of email and social networking sites believe that Facebook, Yahoo, and Google are authorized to engage in controversial and invasive practices implicating user privacy, and that (4) there is presently little reason to expect the development of a robust market for premium privacy-protective email and social networking applications in the United States.

¹ Lior Jacob Strahilevitz is the Sidley Austin Professor of Law at the University of Chicago.

² Matthew B. Kugler is an Assistant Professor at Northwestern University Pritzker School of Law. The authors owe thanks to Omri Ben-Shahar, Yun-chien Chang, David Hoffman, William Hubbard, Randy Picker, Heather Whitney, workshop participants at the University of Chicago Law School, and participants in the Chicago conference on Contracting over Privacy for helpful comments. Adam Woffinden and Taylor Coles provided terrific research assistance. The Russell J. Parsons, Bernard Sang, and Carl S. Lloyd Faculty Research Funds, and the Coase-Sandor Institute for Law & Economics provided generous research support.

1. Introduction

Privacy class actions have become a major financial liability for technology companies. Both Yahoo and Google have been sued over their practices of scanning the contents of users' emails to serve them with personalized advertisements, with plaintiffs alleging that their actions violated the Wiretap Act.³ In each case the potential liability would have been staggering. Plaintiffs would have been entitled to a minimum of \$100 per day of the violation, easily leading to total damages in the trillions for a company with as many users as Google.⁴ Facebook and Shutterfly have been sued for similar privacy violations under the Illinois Biometric Information Privacy Act, 740 ILCS 14 et seq., for their use of biometrics to identify people in uploaded photos. Here, too, liability could be enormous.⁵

In the Google and Yahoo cases, the defenses have turned on the content of the privacy policies active during the relevant period. Yahoo's policy was deemed sufficiently explicit about the email monitoring that a judge ruled that its users had consented to the monitoring, defeating the wiretap claim. That same judge held that Google's policy was not sufficiently clear, however, so its users had not consented. Though the plaintiffs' motion for class certification there was ultimately defeated, limiting the Google's exposure, both the original case and related litigation are still ongoing. (Stempel 2014; Corley v. Google Inc. Complaint, N.D. Cal. Jan. 27, 2016).

The lawsuit against Facebook is still in the early stages, but it appears that it too could turn on whether their users consented to the alleged conduct by agreeing to their privacy policies and whether they were sufficiently informed about how the data would be used. (Pezen v. Facebook Class Action Complaint, N.D. Ill. April 21, 2015; Licata v. Facebook Consolidated Class Action Complaint, N.D. Cal. Aug. 28, 2015). A copycat lawsuit against Shutterfly for its violations of the same Illinois statute, based on similar underlying conduct, has already withstood the defendant's motion to dismiss. (Norberg v. Shutterfly Class Action

³ In re Google Inc. Gmail Litigation 2013; In re Yahoo Mail Litigation 2014; 18 USC § 2511.

⁴ 18 U.S.C. § 2520(2)(B)). Most Gmail users send or receive some email every day, and Gmail has approximately 500 million users worldwide. If Gmail had an average of 50 million American users during the five-year period of alleged violations, then its liability under the lawsuit could \$9 trillion (50 million users * 365 days per year * 5 years * \$100 per user per day).

⁵ The stakes in the Facebook suit are again high because of a \$5,000 minimum statutory damages provision in the Illinois law. (ILCS 14/20(2); Welinder 2012). A back-of-the-envelope calculation reveals that even if no Illinois Facebook user could sue for multiple violations of the law, Facebook's potential exposure is still approximately \$37.5 billion. Approximately 58% of Americans had Facebook accounts as of 2015, and Illinois had about 12.9 million residents at that time. Assuming Illinois residents use Facebook at national average rates, that means there are about 7.5 million Facebook users in the state. Multiplying that figure by \$5,000 yields \$37.5 billion. But if each instance of unauthorized tagging is a separate violation, then Facebook's potential liability could quickly escalate from there. The statutory text seems ambiguous on the question. (740 ILCS 14/20 et seq.)

Complaint, N.D. Ill., June 17, 2015; *Norberg v. Shutterfly Order Denying Defendant’s Motion to Dismiss*, N.D. Ill., Dec. 29, 2015).

In each of these cases, courts have been tasked with interpreting consumer privacy policies. In ruling on Google’s motion to dismiss the initial Wiretap Act suit, the district court assumed that Gmail users read the privacy policies in question and then found that agreeing explicitly with the terms of those policies would not have amounted to consent to the automated email content analysis as a matter of law. As the district court knew, and as scholars have long argued, consumers do not typically read privacy policies and other online disclosures, even for products like Gmail that they use every day. (McDonald and Cranor 2008, Marotta-Wurgler 2011; Schneider and Ben-Shahar, Ayres and Schwartz 2014, Porat and Strahilevitz 2014.) But the “duty to read” is nevertheless very well established in contract doctrine. (Knapp 2015). Courts know that most consumers do not read privacy policies but pretend otherwise for the purposes of contract law and then ask how a reasonable consumer would have interpreted the contract.

Suppose that consumers actually read consumer contracts and privacy policies. What would they understand from them? Would actual consumers draw the same distinctions between, say, the Yahoo and Gmail privacy policies that the district court did? This article addresses that question through an experimental approach, and the results are surprising. After reading actual policy language from Gmail, Yahoo, and Facebook, American users of email and social networking websites largely believe that by using those products they have consented to automated content analysis of their emails and the use of facial recognition biometrics to suggest photograph tags. That is true regardless of whether consumers read versions of those privacy policies (like Yahoo’s) that are extremely explicit, or whether consumers read companies’ older privacy policies, which (at least in the Gmail litigation) a court deemed inadequate to obtain users’ consent. In short, even when consumers do read privacy policies, their beliefs about the nature of their bargains with technology companies seem to depend more on their pre-existing expectations than on the terms of the policies. (Wilkinson-Ryan 2014).

Interestingly, it does not appear that Americans’ views that they have consented to such privacy intrusions stem from normative approval of Google and Facebook’s respective practices. When asked about the intrusiveness of Google and Facebook’s practices, respondents rate these practices as highly intrusive. In light of these reactions, the most plausible interpretation of the data presented here is that email and social networking users believe these practices are part of the bundle associated with Gmail and Facebook, and believe themselves to have accepted that bundle, all the while preferring that the bundle included greater privacy protections.

2. Prior Literature

There is a slowly growing experimental literature on consumer contracts. Some of it, like the present studies, employs random assignment techniques to determine what effects changes in contract language or structure have on consumer behavior. For example, Zev Eigen

randomly assigned online survey participants to conditions that mimic standard contract boilerplate, a compelled choice between two terms, and notice plus choice. (Eigen 2012). He found that respondents assigned to the boilerplate condition were less likely to read contractual terms and also devoted less energy to performing the task the experiment asked them to do. Joshua Mitts randomly assigned a mix of real and fictitious contract terms to respondents and identified surprising / unexpected terms. Such terms were then highlighted with warnings for consumers. (Mitts 2014). He found that the more times warnings about unexpected terms were given to consumers, the less effective each warning was in helping consumers understand the terms of the agreement. And David Hoffman has found that consumers, particularly younger ones, generally regard written contracts as more binding than oral contracts. (Hoffman 2016).

Other experimental research identifies the role that consumer contract language can have in shaping consumers' expectations about the nature of the bargain. Stanislav Mamonov and Raquel Benbunan-Fich found that consumers regard privacy breaches as more disturbing when they are told that the party storing the data had rights to use it than when told that the party storing the data lacked such use rights. (Mamonov and Benbunan-Fich 2015). This research suggests that respondents do care about what's in privacy policies and that such content can shape their understanding of a counterparty's obligations. Other experimental research suggests that attributes like contract length affect consumers' likelihood of accepting or rejecting a written contract. (Plaut and Bartlett 2012). Similarly, the existence of liquidated damages provisions in mortgage contracts affects the extent to which experimental subjects regard contractual breaches as immoral. (Seiler 2016).

A separate literature examines the psychology of consumer contracts. This literature indicates that only a miniscule percentage of consumers read boilerplate contractual language (Marotta-Wurgler 2012), and that parties' expectations about the contents of a contract are driven not only by written terms of the deal but also by moral and legal norms. (Wilkinson-Ryan 2012). We see that the formalization of a contractual arrangement looms large in the lay understanding of what it means to be bound by promises, and contract formation is less of a binary on-off switch than a gradual process where parties feel increasingly bound as the relationship becomes more formalized over time. (Wilkinson-Ryan and Hoffman 2015). Finally, and most relevantly for present purposes, consumers who have signed contracts often feel morally bound to those terms, even when they regard the terms they have agreed to as substantively unfair and when their agreement to those terms causes them to suffer harm. (Wilkinson-Ryan 2014).

Another relevant experimental literature explores the existence of a "privacy paradox." Privacy paradoxes arise because Americans often say they care a great deal about privacy and yet they are willing to permit third parties to obtain sensitive information about them in exchange for relatively inexpensive goods and services, or in exchange for longshot odds to win a prize in a random drawing (Acquisti 2010, Holland 2010, Swire 1999). The diminished value placed on privacy may stem in part from framing effects. (Acquisti, John, and Loewenstein 2013).

3. Data and Empirical Approach

3.1 The Sample

Toluna, a professional survey research firm with an established panel, administered a survey to a weighted sample of 1,441 adult US citizens between May 26, 2015, and June 2, 2015. Data from some of these was discarded because of abnormally fast survey completion times and failed attention checks, leaving a final sample of 1382. The median age of respondents was 47 (range 18-89, mean: 46.62, SD = 16.37). Females comprised 49.8% of the sample. Compared with the population in the US census, a higher percentage of the panel had completed high school or at least some college coursework, but the educational attainment of the sample was otherwise similar to that of the adult census population. 79.9% of the sample self-identified as White, 13.0% as Black, and 4.1% as South or East Asian. On a separate question, 16.2% of the sample reported that they are Latino or Hispanic. Respondents were asked their political orientation on a scale of 1 (very liberal) to 7 (very conservative), with a mean response of 4.16 (SD = 1.78), indicating an ideologically moderate sample. The Gmail and Facebook questions were administered at the end of a 10-15 minute survey that included questions for other papers on topics such as Fourth Amendment privacy expectations and trademark questions designed to assess attributions of product sponsorship.⁶

Participants were screened on the basis of whether they reported having email accounts for the Gmail questions and whether they said they had Facebook accounts for the Facebook questions. That left 1377 potential respondents to the email questions and 1,052 potential respondents for the Facebook questions.⁷ Approximately 76.1% of the sample were therefore Facebook users. This utilization rate is close to the one produced by a Pew Research study conducted a few months earlier, which found that 72% of American adults with Internet access used Facebook.⁸ In each instance, eligible respondents were randomly assigned one of three

⁶ These survey results are discussed in Kugler & Strahilevitz 2016, and Kugler 2015, respectively.

⁷ Facebook users were, on average, slightly younger than non-Facebook users (Users mean = 45.06, SD = 16.17; Non mean = 51.58, SD = 16.09). The Facebook user population was also more female (52.4%) than the general sample. The racial breakdown was roughly equivalent, however (79.0% White, 13.6% Black, 4.0% South or East Asian). Note that 28 respondents indicated that they had Facebook accounts but did not answer any of the other Facebook-related questions, so they were dropped from this experiment.

⁸ (Duggan 2015). The Pew study reports that 62% of all US adults are Facebook users. Although our Toluna sample is census-weighted, Americans without Internet access were necessarily excluded from the online survey. This exclusion does not seem problematic given our interest in learning how consumers of privacy policies and online apps understand those policies. The exclusion of those without Internet access (15-16% of the adult population) largely explains the disparity in education levels between our sample and the adult population. (Perrin & Duggan 2015). The (declining) American population of non-Internet users is older, lower income, less educated, and more rural than the population of Internet users. (Anderson & Perrin 2015).

“privacy policies” for both the Gmail and Facebook questions. In each instance the privacy policy language subjects read was taken from actual language that Google or Facebook employed at some point in time.⁹ The policy language varied in terms of how explicit it was about Google and Facebook’s data practices. Not surprisingly, the current policy language (posted after the main lawsuits at issue here were filed) is more explicit about company practices than the pre-lawsuit language.

3.2 The Survey Instrument

The randomization strategy in the experiment allows for a clean test about what effect differing policy language has on consumers’ views of what they have agreed to. The difference in the new language and old language was (to these lawyers’ eyes, at least) dramatic enough to warrant the following pre-experiment hypothesis: Lay understandings of privacy policies would depend significantly on the policy language chosen. Given the prominent display of just the relevant language to respondents, enough consumers would read the privacy policy excerpts closely to render the substantial differences between the old and new privacy policies significant.

All respondents were asked to assume that when they signed up for email they agreed to permit advertisements to be shown next to their inboxes in exchange for a free account, and they were also asked to assume that they had read the terms and conditions when signing up for the account. They were then shown randomly assigned privacy policy language that concerned whether these advertisements could be personalized. For example, some saw Gmail’s current language, which is quite explicit: Email provider’s “automated systems analyze your content (including emails) to provide you personally relevant product features, such as customized search results, tailored advertising, and spam and malware detection.” Others saw much vaguer language that Gmail used to post: [Email provider] “reserves the right to pre-screen, review, flag, filter, modify, refuse or remove any or all content from any service. For some services, [email provider] may provide tools to filter out explicit sexual content.” Gmail argued unsuccessfully in court that their users’ agreement to even that very vague language granted Google’s consent to show personalized advertisements to Gmail users.

Respondents were also asked other questions targeted at issues beyond whether they had consented to the legally relevant conduct by Google and Facebook. Respondents to the Gmail survey were asked “On a scale of 1 to 10, how intrusive is the email provider’s automated email scanning and ad personalization practice?” After answering this question, they were asked: “If there were an option to keep the same email account but pay some amount of money to avoid having the automated systems analyze email content for the purposes of showing you personalized advertisements, how would you respond? (1) I would keep the free email account with the automated email analysis and personalized advertisements. (2) I would be willing to pay some amount of money to avoid the automated analysis.” Respondents who

⁹ The Gmail questions used both Google’s current language from 2015 and the circa 2011 language quoted in the Gmail litigation. The Facebook questions used Facebook’s current language and earlier versions of related privacy policies obtained via the Internet Archive Wayback Machine.

selected option 2 were asked how much money they would be willing to pay per year for a more privacy protective email product.

Respondents to the Facebook question, all of whom had Facebook accounts, were randomly shown various Facebook privacy policy language and then asked four questions designed to elicit responses that would shed light on whether Facebook had complied with its obligations under Illinois law. Again, as can be seen from perusing the policy language in the Appendix, differences in the privacy policies seemed stark at first blush. All Facebook respondents were then asked: “Did Facebook’s language (above) inform you that information about your facial features was being collected and stored?” “Did Facebook’s language (above) inform you of the reason why information about your facial features was being collected, stored, and used?” “Did Facebook’s language (above) inform you of the length of time for which information about your facial features would be stored?” (740 ILCS 14 et seq.) And then finally, they were asked the consent question: “Would your decision not to adjust your Timeline and Tagging settings allow Facebook to collect, store, and use information about your facial features?”

Respondents to the Facebook questions were then asked to rate on a scale of 1-10 the intrusiveness of Facebook’s use of facial recognition software to suggest tags for people whose faces appear in uploaded photos.

4. Results

Given the substantial differences between the privacy policy language that email and social networking site users were shown – all this language is reproduced in the Appendix – we predicted that our respondents who saw the highly explicit disclosures from Google and Facebook would be more likely to say that their decision to leave their privacy preferences unchanged after reading the relevant privacy policies allowed Google and Facebook to engage in the content analysis and facial recognition practices at issue. Surprisingly, the data did not support that hypothesis. Regardless of what language respondents were shown, they had statistically indistinguishable views about what practices their inertia would have authorized.

4.1 Experiment 1: Gmail Results

In the Gmail experiment, random assignment to one of three conditions – Google’s very explicit current privacy policy, Google’s moderately explicit historic section 17 language, or its least explicit historic section 8 language – had no significant effect on consumers’ judgment about what they had authorized Google to do to their emails. Nor did the privacy policy language have any significant effect on the perceived intrusiveness of Google’s automated content analysis of their customers’ emails. Differences in language that lawyers and judges would deem critical made no evident difference to a representative sample of adult American email users. (Compare Reidenberg et al. 2014). The differences in means are even in the wrong direction for the comparison between the most explicit policy and the moderately explicit.

Table 1: Responses to Gmail Consent Question by Privacy Policy -- “Would your agreement to this provision allow the email provider to direct its automated systems to scan the contents of the emails you send and receive and show you personalized advertisements?”

	Current Language (most explicit)	Section 17 (moderately explicit)	Section 8 (least explicit)	Overall (ignoring condition)
Definitely Allowed	28.1%	28.1%	23.5%	26.6%
Probably Allowed	35.7%	40.2%	39.6%	38.5%
Probably Not Allowed	13.9%	10.6%	15.6%	13.4%
Definitely Not Allowed	22.2%	21.2%	21.3%	21.3%
Mean (Allowed=1, Not = 4)	2.30 (1.10)	2.25 (1.08)	2.35 (1.06)	2.30 (1.08)

The frequency differences across condition are not significant. $\chi^2(2, N = 1363) = 8.38, p = .21$. Numbers in parentheses are standard deviations and means do not differ across conditions.

Moreover, in every condition, most respondents say that if they read the short privacy language at issue and then did not change their privacy settings to prohibit content analysis, Google would be authorized to engage in the automated content analysis. Roughly two-thirds of the sample expressed this view in all three conditions.

One possible interpretation of this result is that email users like receiving personalized advertisements and do not mind the automated content analysis of their email that facilitates this personalization. On this interpretation of the data in Table 1, consumers’ normative views would be driving their answers to the question of what Google can do. But this interpretation is not supported by the intrusiveness data, shown in Table 2.

Table 2: Responses to “On a scale of 1 to 10, how intrusive is the email provider’s automated email scanning and ad personalization practice?”

Email Condition	Intrusiveness Mean	Std. Dev.	N
Most Explicit	7.60	2.47	445
Moderately Explicit	7.62	2.34	463
Least Explicit	7.65	2.43	455
Total	7.63	2.41	1363

Dep. Variable – intrusiveness: $F(2, 1360) = .06, p = .95 \eta^2 = .000$.

Mean intrusiveness responses for Google’s conduct is 7.63 on a 10-point scale. Consumers are saying that they regard the automated content analysis as rather creepy, but nevertheless authorized, even when presented with language that few lawyers would regard as consenting to the practice at issue. Intrusiveness ratings, predictably, were not significantly affected by whether respondents saw more explicit or less explicit privacy policies. Those who believe

Google is less authorized to scan emails view the practice as slightly more intrusive ($r(1363) = .192, p < .001$), but the effect size is very small.

Our results are consistent with the privacy paradox as well. Although the mean respondent rated automated content analysis of emails as a 7.63 out of 10 on an intrusiveness scale, just 35.4% of the sample expressed a willingness to pay any amount of money to receive a version of their email service that did not use automated email content analysis to deliver personalized ads. Among the roughly one-third of the sample that was willing to pay some amount of money, the median willingness to pay was \$15 per year. Just 3% of the sample expressed a willingness to pay more than \$120 per year for such an email service.¹⁰

Perhaps this data indicates that the intrusiveness ratings offered by our respondents are not to be taken seriously. Maybe the 7.63 intrusiveness figure is just cheap talk. On this reading of the data automated email content analysis is not a serious concern for most Americans, which explains why they feel that Google is allowed to engage in the practice even without explicit ex ante warnings. Another possibility is that users of the Internet have grown accustomed to “free” email, news, weather, media content, etc., such that putting email behind a paywall prompts significant resistance even when doing so would create a substantially more privacy-protective product. (Dou 2004; Acquisti, John and Loewenstein 2013). Alternatively, perhaps consumers say they are reluctant to pay any dollar amount for a privacy-protective email account precisely because they know that other email services (e.g., Hotmail) exist and respondents correctly surmise they do not engage in automated content analysis.¹¹ Finally, it may be that by purchasing a right to be free of automated content analysis, consumers would be acquiring just a tiny privacy enhancement that would make little difference given other invasive practices. Perhaps if consumers could bundle together a lack of content analysis with limits on behavioral marketing, the commercial use of geolocation, facial recognition software, and sharing of personal information across websites, they would be willing to fork over a more meaningful amount of money.

In any event, the shortage of consumers willing to pay meaningful sums for more privacy-protective email services suggests there may be a limited market for premium products that protect user privacy. Recent estimates indicate that a year’s worth of data is worth \$50 to \$5000 per consumer to Google and \$45 to \$190 per consumer to Facebook. (Howe 2015). The sorts of fees they’d be able to obtain from users for greater privacy-protection are relatively small potatoes, though it is conceivable that enhanced data security would prompt a more robust response from consumers. In any event, the shortage of consumers willing to pay

¹⁰ Although Hoffman (2016) finds significant differences in the ways that younger respondents and older respondents think about contracting online, we found significant age effects for neither willingness to pay nor the perceived intrusiveness of Facebook and Google’s conduct.

¹¹ Microsoft’s privacy policy states “we do not use what you say in email, chat, video calls or voice mail, or your documents, photos or other personal files to target ads to you.” See <https://privacy.microsoft.com/en-us/privacystatement> (last updated January 2016).

meaningful amounts for more privacy protective email accounts plausibly explains why Google has not offered privacy-differentiated email products.

Our data provide information that permits some inferences to be drawn about the dynamics at play. It does not appear that differential views about the intrusiveness of automated email content analysis are driving users to one email provider or another. Mean intrusiveness ratings were not significantly different among Gmail, Yahoo, AOL, and Hotmail users. $F(3, 1136) = 1.44, p = .23 \eta^2 = .004$. Nor do consumers appear to be choosing their email providers based on their privacy preferences and company policies more broadly. When we analyzed responses to questions about the intrusiveness of Facebook's facial recognition software (discussed below) based on what email providers respondents use, there were no significant differences.

It is less clear if awareness of different company practices affect respondents' assessments of whether automated content analysis is permitted. Gmail users were significantly more likely than AOL users to believe that email content analysis was permitted, but so were Hotmail users, and the effect sizes were small in any event.¹² (AOL and Hotmail evidently do not perform automated content analysis on their customers' emails.) Given that respondents were asked about whether their own email providers were allowed to engage in automated content analysis, it seems that at most a small portion of the population is attentive to the differences between Google's content-analysis and Hotmail's lack thereof.

Our study also generated mixed evidence on the question of willingness to pay. On the one hand, respondents willing to pay some amount of money to avoid content analysis rated the intrusiveness of the content analysis at 8.65 (1.83), whereas those unwilling to pay any amount rated it at 7.06 (2.50).¹³ On the other hand, the amount people were willing to pay (above zero) bore no relationship to either the perceived intrusiveness or the authorization of automated content analysis.

4.2 Experiment 2: Facebook Results

Under the Illinois Biometric Information Privacy Act, the pertinent legal questions are whether Facebook informed its users about (a) the fact that information about their facial features was being collected and stored, (b) the reason why information about their facial features was being collected, and (c) the length of time for which information about their facial features would be stored. In addition, the law renders germane the question of (d) whether Facebook had its users' permission to collect and store information about their facial features. Each of these four questions depends on Facebook users' understanding of Facebook's terms of service. This experiment was designed to test whether, if Facebook users had read the relevant

¹² Gmail 2.23 (1.07); Yahoo 2.39 (1.09); AOL 2.45 (1.06); Hotmail 2.19 (1.03); Total 2.31 (1.07); $F(3, 1136) = 2.93, p = .033 \eta^2 = .008$. Gmail and Hotmail both were significantly lower than AOL and Yahoo ($p < .05$), and did not differ significantly from each other.

¹³ $F(1, 1136) = 149.49, p < .001 \eta^2 = .10$.

information, they would feel that Facebook had adequately informed them of its practices and obtained their authorization to engage in them.

There was a clear consensus among respondents on all four questions, and the consensus is particularly interesting on the third of the four questions.

Table 3: Responses to Facebook Questions: Percentage of Respondents Answering “Yes”

	Policy – “We collect” (Least explicit on Facebook’s actions and purposes)	Policy – “We use” (Least explicit on Facebook’s actions, more disclosure on purposes)	Policy – “When Someone Uploads” (Most explicit on Facebook’s actions, less on purposes)	Total
Did Facebook inform you about collection and storage?	67.8%	65.9%	70.7%	68.2%
Did Facebook inform you of reason for collection, use, and storage?	59.0%	61.8%	67.1%	62.5%
Did Facebook inform you about length of time information would be stored?	35.1%	34.0%	30.7%	33.3%
Does leaving settings unchanged allow Facebook to collect, use, and store information?	63.5%	62.3%	61.0%	62.3%

Despite a sample size of 1052 respondents, on none of the questions presented in Table 3 is the language from Facebook’s privacy policies having any significant effect. More than two-thirds of the sample regards themselves as having been informed of Facebook’s collection and storage of their biometric information after having read any of Facebook’s current or historic policy language. And an only slightly lower percentage of Facebook users view Facebook’s language as informing them of the *purpose* of Facebook’s use and collection. Again, the wording of the policy language at issue made no significant difference, even though in one condition the language was very explicit about the purposes of Facebook’s collection of information and in the other it was not. Similarly high percentages of respondents said user inaction with respect to privacy settings authorized Facebook’s facial recognition practices.

Viewed in context, the most striking set of Table 3 responses are those to question three, which asks about the length of time for which Facebook is retaining its information. In

none of these conditions did the privacy policy language provided to respondents address the duration of storage explicitly. Up to 67% of the respondents seem to have noticed this. This reversal of the usual ratios across all three conditions suggests several possible implications. First, it seems that at the very least a third of the sample is reading the lengthy privacy policy language in the prompt carefully. These are the respondents who flip from a pro-Facebook stance on the other questions to an anti-Facebook stance on question 3. Second, it is possible (though unlikely) that whereas Facebook users have intuitions about the fact that a facial recognition algorithm is being used and the reason why it is being used (perhaps based on their use of the feature on Facebook), they lack a strong prior about the length of time for which facial recognition information should be retained, so the privacy policy language may play a larger role than context in shaping their understanding.¹⁴ Third, unless there is other privacy policy language that Facebook can cite,¹⁵ it appears plausible that, according to consumers, Facebook's facial recognition feature has been violating one provision – though only one provision – of the Illinois law. To confirm this hypothesis we would need to test the effects of Facebook's data retention duration language on consumers.

Respondents were also asked about the intrusiveness of Facebook's practice of using facial recognition software to suggest tags for people whose faces appear in uploaded photos. Mean responses were a little lower than in the Gmail question (mean = 7.29 out of 10, SD = 2.36) and did not differ by condition ($F(2, 1044) = 1.30, p = .27, \eta^2 = .002$). Thus, it does not appear that exposure to different policy language affected consumers' underlying beliefs about how problematic Facebook's practices are. Once again, majorities of consumers appear to regard Facebook's practice as troubling yet authorized. Comparing across conditions between authorization and perceived intrusiveness responses did not yield significant results. ($p = .396$).

5. Replication

Whenever a null result is observed in this type of vignette experiment, one possibility is that participants simply did not attend to the materials. As we know from prior research, boilerplate policy language may encourage consumers to tune out fine details. (Eigen 2012). We

¹⁴ Research from Pew suggests that Americans generally do have articulated priors about the length of time for which their personal data should be retained. Just 4% of respondents said that social media or online video sites should be able to retain their data for "as long as they need it." (Rainie 2016).

¹⁵ As of January 14, 2016, Facebook's Data Policy provided in pertinent part: "We store data for as long as it is necessary to provide products and services to you and others" and equivocated on how much data would be eliminated if the account was deleted. We did not show this clause to our experimental subjects because this language is plausibly too vague and indefinite about the duration of data retention to satisfy the Illinois statute, and the permanent retention of biometric information gleaned from photos uploaded by other users who do not delete their accounts could well violate the Illinois statute. More broadly, the failure to disclose the duration of data retention appears to be quite commonplace in the United States. (Marotta-Wurgler 2017).

therefore conducted a replication study that aimed at assessing how carefully participants read the provided materials and whether more attentive participants differed from less attentive ones. The new features in this replication were: 1) a measure of how long participants spent on the main Facebook and email scenario pages; 2) a manipulation for half of the participants in the email portion of the study that asked them to explain why they thought monitoring was or was not allowed (to encourage deeper thought); 3) a manipulation within the Facebook portion of the study that either gave or omitted information about how long the information would be retained; 4) self-report questions on both the Facebook and email scenarios asking participants how well they felt they understood the materials (10-point scale from 1 to 10); and 5) the imposition of a more cognitively demanding attention check that permits us to test our first study's results on a subsample of our most attentive respondents.

The replication experiment also introduced a new email condition that included privacy policy language from Yahoo. The same judge who had held the less-explicit Gmail language to be inadequate was satisfied by this alternative. Adding a condition with this language therefore addresses any concerns about our determination that Gmail's current privacy policy – unlike the earlier policy language that is at issue in the Gmail litigation – would be deemed sufficiently clear and precise to secure consent from consumers who read it.

The procedure was otherwise as before, though the Facebook and email questions came much earlier in the survey, immediately after the background questions, and an effort was made to recruit more participants who had not completed high school or attempted college course work to have a more representative mix of education. A total of 1300 participants were recruited, 1283 of whom completed the email questions and 1045 of whom completed the Facebook questions.¹⁶

In general, the main results of the first study replicated, participants appeared to be paying attention, and more attentive participants did not draw greater distinctions between scenarios than less attentive ones. Three different email scenarios were used: the least explicit from Study 1, the most explicit, and a version from Yahoo's own terms of service that was even more explicit than any employed by Gmail "automated systems scan and analyze all incoming and outgoing communications ... to match and serve targeted advertising...". Results again showed no significant differences in whether the policies allowed the described monitoring, or on perceived intrusiveness. There was a slight difference across condition on whether the participants felt they understood the policy: participants were slightly less confident that they understood the least explicit Gmail policy than either of the other two ($p < .06$).

¹⁶ The median age of respondents was 44 (range 18-90, mean: 45.72, SD = 16.08). Females comprised 51.0% of the sample. 81.5% of the sample self-identified as White, 10.5% as Black, and 3.5% as South or East Asian. On a separate question, 15.5% of the sample reported that they are Latino or Hispanic. Respondents were asked their political orientation on a scale of 1 (very liberal) to 7 (very conservative), with a mean response of 4.11 (SD = 1.75), 11.38% of the sample had not finished high school, 30.38% had high school diplomas, 29.08% had some college experience, 19.00% had college degrees, and 10.15% had some kind of graduate degree.

Table 4: Mean responses and significance tests for the email questions from the replication study.

	Gmail Policy, Least Explicit		Gmail, Most Explicit		Yahoo, Like Gmail Most		Total		F	<i>p</i>
Allow Monitoring	2.39	(1.10)	2.34	(1.12)	2.26	(1.12)	2.33	(1.11)	1.68	.19
Intrusive	7.37	(2.39)	7.57	(2.31)	7.70	(2.22)	7.55	(2.31)	2.33	.10
Understand Policy	7.74	(2.07)	8.01	(2.01)	8.08	(1.99)	7.95	(2.03)	3.34	.04

Further, asking participants to explain why they thought the policies they were given did or did not allow monitoring had no effect on whether they thought the policies permit such monitoring ($F = 1.37$) and did not interact with condition to predict whether permission was imputed ($F = 0.47$). In fact, the only effect of requiring explanations was to make participants take longer on the screen ($F = 123.09, p < .001$).¹⁷

A series of linear regressions was conducted attempting to predict the score on the Allow Monitoring dependent measure from condition and its interactions with either self-reported understanding or time spent on page. Neither variable interacted with condition, indicating that people who spent longer with one policy or another, or felt they better understood one policy or another, did not differ from other participants in whether they thought the policy allowed monitoring.¹⁸

For the Facebook scenarios, the most (“When Someone Uploads”) and least (“We Collect”) explicit policies reprised their roles from Study 1, but a new version of the most explicit was created that included the line “We automatically delete all facial recognition information once it has been stored in our system for three years.” This changes the correct answer to the length of time the information is stored. As can be seen below, participants are sensitive to this change: the majority of the sample in that condition recognize that they’ve received this information, significantly more than in the other conditions. That said, about 37% of the sample answered this question incorrectly, indicating that they did not read the policy closely or that their priors overwhelmed the policy language. The other questions, assessing what the policy actually means for users, do not produce different answers across conditions, replicating Study 1. Perceived understanding also did not differ across condition ($F = .67$, overall mean = 7.44, SD = 2.26)

¹⁷ Since time spent on the page was not normally distributed (some participants were on the page for a long time), the variable was capped at 250 seconds for this and all subsequent analyses.

¹⁸ Interestingly, there were two main effects. Those who spent longer on each page (regardless of condition) were less likely to say they believed the policies allowed the monitoring (standardized beta = .116, $p < .001$), and those who felt they better understood the policies were more likely to believe monitoring was allowed (standardized beta = -.149, $p < .001$).

Table 5: Facebook Question Replication Responses: Percentage of Respondents Answering “Yes”

	“When Someone Uploads”	“When Someone Uploads” (with 3 year limit)	“We collect”	Chi Square	<i>p</i>
Did Facebook inform you about collection and storage?	73.68%	77.30%	71.74%	2.81	0.24
Did Facebook inform you of reason for collection, use, and storage?	62.57%	61.10%	65.73%	1.59	0.45
Did Facebook inform you about length of time information would be stored?	39.59%	62.64%	34.17%	62.45	< .001
Does leaving settings unchanged allow Facebook to collect, use, and store information?	64.12%	59.83%	65.63%	2.62	0.27

Several regressions were conducted to see whether the effects of perceived understanding or length of time on the page affected responses to these questions differently depending on condition. For the understanding question, there were no significant interactions, meaning that those who thought they understood the prompt well did not come to different answers depending on which prompt they read.¹⁹ For time spent on page, the only interactions were on the duration of data retention question.²⁰ Those who spent longer on the page were more likely to come to the “right” answer on that question, meaning “yes” for the duration limit condition and “no” for the other two conditions.

Finally, the introduction of our new 3 year time-limit condition in the replication experiment permitted us to apply a relatively demanding new attention check to our sample. After our first experiment, we hypothesized that respondents’ assessment of whether email content analysis was permitted would not differ across conditions even among our most

¹⁹ There were significant main effects of self-reported understanding on the first three questions, odds ratios of .53; .69; and .66 respectively (all *ps* < .001).

²⁰ This effect is easier understood in terms of an ANOVA. There was a significant interaction between Facebook condition and the answer to the duration limit question on time spent on page. $F(1, 998) = 23.51, p < .001, \eta^2 = .05$. “When Someone Uploads” Yes = 46.99 (49.89); No = 100.75 (66.09); With 3 year limit Yes = 95.11 (69.09); No = 80.72 (60.26); “We collect” Yes = 70.94 (76.72); No = 112.67 (74.30). In the duration limited condition, those saying yes took significantly longer. In the others, where this was the wrong answer, those saying yes took significantly less time (*ps* < .05). The binary logistic regression version of this analysis is available from the authors.

attentive readers. In the replication study we tested this hypothesis by examining whether our headline results would change if we omitted the Yahoo and Gmail responses of all subjects who answered the Facebook data retention questions incorrectly. Even respondents who read the Facebook questions closely enough to notice the presence or absence of a single sentence buried in a paragraph from a privacy policy did not differentiate between Yahoo's legally adequate and Gmail's legally inadequate privacy policies in terms of whether content analysis was authorized. ($F < 1$).

6. Discussion

The key lesson from both the Facebook and email data is that users of email and social networking sites appear to regard even highly ambiguous privacy policy language as authorizing controversial company practices that implicate their personal privacy. Tess Wilkinson-Ryan finds a similar result in the context of other boilerplate consumer contracts. (Wilkinson-Ryan 2014). Though federal courts determined that Yahoo's privacy policy informed email users of the company's automated content analysis and that Gmail's privacy policy did not, American email users did not differentiate between the purportedly adequate and inadequate policies. To the contrary, they thought that agreeing to either policy would establish their consent to automated content analysis.

What explains the divergence between lawyerly judgments and lay consumers' judgments about what constitutes consent? One possible explanation is that consumers had formed strong priors about the sort of privacy-related conduct that companies are permitted to engage in, and these priors inform their understanding about what they agree to when they use Gmail or Facebook without changing their privacy settings. (Martin 2015). Even when consumers are familiar with the formal law and written policy language, expectations are also driven by social norms. When consumers interpret contracts, they bring in these priors and integrate their beliefs with the policy language to produce an understanding of the bargain to which they are agreeing. (Hoffman and Wilkinson-Ryan 2012). Consumers may not like the bargain in all material respects – and their intrusiveness scores suggest discomfort with automated email content analysis and the automated use of facial recognition software – but they seem to believe the privacy sacrifices inherent in their use of email and social networking sites outweighs those costs.

When faced with data like this and a consent defense by a defendant who invokes this sort of empirical evidence, what should a court do? In our view, data such as this, collected using rigorous survey techniques and analyzed by academics with no skin in the game, ought to play a large role in litigation over privacy policies in particular and consumer contracts in general.²¹ Under such an approach, consumer contract interpretation would become a question of fact rather than a question of law. Where a consensus emerges among consumers as to the contours of a deal, this consensus understanding would become the contract's meaning, even

²¹ Ben-Shahar and Strahilevitz have developed this argument in much more depth in a working paper, tentatively titled A New Approach to Contract Interpretation: The Consumer Survey Method.

among those consumers who had subjective views of the contract that placed them in the minority. This survey-driven approach would represent a break with American law's dominant paradigm for contract interpretation.²² Under a survey-driven approach the interpretation of consumer contracts would more closely resemble what courts do in trademark litigation, where consumer surveys are dispositive (Diamond and Franklyn 2014).

Though at first blush this change in the law would make the law more hostile to business interests (by making it harder to win a motion to dismiss in a contract suit) and friendlier to plaintiffs' interests, this result is hardly inevitable. Battling over the legal meaning of contractual terms is not cheap. A consumer-survey-driven approach to contract interpretation would resolve cases at a later stage, but that does not mean that more money would be spent before resolution. Legal research that takes place early in litigation now could be replaced with survey research. To the extent that dominant survey methodologies emerged quickly, then the parties could promptly settle in the shadow of their respective experts' survey results. Indeed, a lot of current contract claims might never be brought in the first instance because plaintiffs' attorneys would have a relatively inexpensive way to test whether a breach of consumer contracts claim would be viable. It is plausible, though by no means certain, that prompting the law to focus on ordinary consumers' actual understandings of contractual provisions would be more efficient than the current approach.

The goal of companies designing privacy policies and consumer contract language should be to inform consumers about what the companies are doing and why they are doing it. Companies already field test their products extensively. For similar reasons, they should field-test their policy language on consumers and avoid presuming that the only information consumers have is what is disclosed in the policy language. It's precisely because lawyers are trying to cram so much information into policies that policies become unduly lengthy, and the result is they go unread entirely by rational consumers. (Ben-Shahar and Schneider 2014). The meaning of a consumer contract is a product of consumers' expectations and the contract language, with the former seemingly looming larger than the latter in some contexts. The product is readily measurable, even if teasing out what work the expectations are doing and what work the language is doing is more complex. At least in the instance of Gmail, privacy policy language chosen by Google and the other information that consumers are receiving or intuiting from various sources does adequately inform most consumers about the nature of the bargain.²³

²² See, for example, *Antilles Steamship Co. Ltd. v Members of American Hull Ins Synd*, 733 F2d 195, 204 n2 (2d Cir 1984) (Newman concurring). One paper comes close to advocating such an approach, but uses the approach to resolve a hypothetical question about the likelihood of scarce goods being available in the future, as opposed to a question of what the contract language itself means. (Olazabel, Marmorstein and Sarel 2014).

²³ If this approach were adopted, it is possible that there would be certain contract provisions that survey respondents would find so surprising / unbelievable that firms could never successfully integrate them into a bargain, no matter how explicit the contractual language employed. We think this data-

Several important caveats remain. First, we know that both (a) consumers very rarely read privacy policies and (b) courts adjudicating class action cases nearly always impose a duty to read on consumers. There may be sensible reasons for the courts to proceed on that basis, particularly at the motion-to-dismiss stage or the summary judgment stage. But if they do assume that consumers read these contracts, it seems highly problematic to assume an interpretation of those contracts that relatively few lay readers of those contracts would share. The duty to read can't possibly mean a "duty to hire a lawyer to read in a lawyerly way." Can it?

Second, in assessing the generalizability of these results, it is important to recall that our respondents were only asked to read a short excerpt of a much lengthier privacy policy. Respondents were not charged with scanning a dense policy and finding the relevant provision. Had we asked respondents to read a lengthier policy carefully, few would have been incentivized to comply. On the other hand, the Gmail and Facebook questions in our first experiment were presented to our respondents toward the end of a 10-15 minute online survey that also asked them a number of questions about Fourth Amendment privacy questions and trademark issues. That was in part the rationale for our replication study, which placed the privacy questions much earlier in the survey. In any event, the results here should be conceived of as relevant to the question of "what would happen if consumers actually read the pertinent parts of privacy policies?" an inquiry that, though hypothetical, winds up being outcome-determinative in a great many litigated cases.

Third, there is an adaptive preferences problem built into our survey methodology that could affect the interpretation of the results. The Facebook experiment was limited to respondents from a nationally representative sample who said they have Facebook accounts. The respondents therefore had already been exposed to Facebook's tagging suggestions, and many may have already realized that Facebook employed facial recognition software to suggest tags. This previous exposure had benefits and drawbacks. One benefit is that many consumers already understood a technological feature that might have been difficult to explain otherwise. (For reasons related to the complexity of the technology we did not ask non-Facebook users to answer the experiment's questions.) But a drawback is that by the time Facebook was sued and we presented respondents with our survey, Facebook had been employing facial recognition technology for nearly five years. (Ducklin 2010). Facebook users' initial understanding of Facebook's practices is arguably as relevant as Facebook users' contemporary understanding of Facebook's practices. The problem is present too in the Gmail survey, where the firm's practice was again longstanding by the time the survey launched. To be sure, the lack of large differences in the responses of Gmail users and demographically similar Hotmail users alleviates some concerns about conditioned responses. Still, as a result of these issues, our study lacks a clear "before" to go with its "after" result. Because it takes time to get a survey developed, approved, funded and launched, it is unlikely that third party researchers will ever be able to test consumers' understandings of companies' new practices before those practices have been implemented. But firms themselves might hire reputable academic researchers to obtain data

driven approach to contract unconscionability might be more appealing than existing approaches. In any event, in this study we have not identified any such terms.

that predates consumer adaptation to a new feature. That said, in both the *Facebook* and *Gmail* litigation, plaintiffs are seeking continuing damages over a period of several years. Even if we cannot identify precise consumer sentiment at the time a controversial practice began, understanding contemporary responses may help place an upward bound on the damages that are appropriate in any given case.

Finally, there is a hard question of what to do with respondent heterogeneity. When presented with language that (to our lawyer eyes anyway) very clearly informs Facebook users of the reasons why Facebook is collecting facial recognition data, 38% of our respondents said that Facebook did not inform them of the reasons for the data collection. A similar percentage of respondents (between 34% and 40%) in the initial and replication surveys provided an objectively incorrect answer to the question of whether Facebook had informed its users about the length of time for which it would be retaining biometric information. And when presented with language that (again, in our judgment) unambiguously informs readers that facial recognition of software is being used to collect data used for suggesting photo tags, 29% of our respondents said the language failed to do so. With any survey instrument, there are going to be some people who do not read very carefully or answer most questions at random but nevertheless answer standard attention check questions correctly, and there will be others who have sufficiently strong views about the facts or morality of an issue to not be swayed by any exculpatory contract language. It appears that in our experiment, those groups combined to form somewhere between 25 and 40 percent of the overall sample. In a world where lawyers have determined that contract or policy language should have some efficacy in shaping consumer expectations, the fact that 30% or 35% of a sample articulates the view that particular policy language with which they were presented is inadequate should not sway a court unduly.

As one examines pleadings in cases like *In re Google Inc. Gmail Litigation*, *In re Yahoo Mail Litigation*, and *In re Facebook Biometric Information Privacy Litigation*, the absence of empirics about how consumers respond to terms of service language is striking. This is information that litigants (or better yet, social scientists) ought to be producing and that courts ought to be evaluating. (Martin 2014). The survey results presented here were neither particularly difficult nor costly to gather. The total costs for our first survey sample were \$4550, but this sample was used to provide the data for this project as well as three other research papers dealing with disparate topics. Compared to a few billable hours of a good lawyer's time, such experimental research is a bargain. And for a corporation that is trying to limit its exposure to class action suits, making nationally representative consumer survey results legally dispositive could be a blessing. Instead of engaging in guesswork about which boilerplate language courts would regard as adequate or inadequate for the purposes of securing consumer consent, corporations could make an ex ante determination that is presumably likely to remain stable down the road. (Kugler & Strahilevitz 2017).

7. Conclusion

It is well-understood that consumers typically do not read boilerplate privacy policies and that, for the purposes of determining whether consumers have consented to particular companies' privacy practices, courts nevertheless assume that consumers did read those policies. Our experiments suggest that even if a large number of consumers did read controversial privacy policies, their interpretations of those policies and of what conduct they had authorized would differ from conventional legal interpretations of those policies' meaning. More precisely, consumers seem to regard themselves as having authorized several controversial privacy-related practices by Google, Yahoo, and Facebook regardless of whether they were randomly assigned to read vague language that doesn't seem to explain the corporate practices in any meaningful detail or precise language that describes the corporate practices at issue with admirable clarity and specificity.

These experimental findings suggest that differences in policy language that are quite salient to lawyers are essentially irrelevant to consumers. Context, experience, and norms, rather than privacy policy language, seem to benchmark consumers' understandings about what conduct they are authorizing, and that is the case even in those instances where one can be confident that consumers have read the relevant policy language rather carefully. Moreover, the experiments reported herein suggest that normative priors about what corporate practices are more or less invasive do not significantly affect most consumers' understandings about what companies like Facebook, Yahoo, and Google are authorized to do. Even though consumers think that uses of facial image recognition and automated email content analysis are invasive, they still regard even vague and imprecise policy language as authorizing Facebook, Yahoo, and Google to engage in those practices. Finally, our experiments provide significant reason to doubt that market forces will significantly incentivize firms to offer privacy-protective alternatives to free services that enhance email privacy. Although consumers dislike automated content analysis, their willingness to pay for a version of Gmail that does not perform content analysis is quite limited, and there is no evidence to indicate that concerns about email content analysis are presently driving consumers to choose substitute email services that eschew email content analysis.

References

Alessandro Acquisti, Leslie K. John, and George Loewenstein, What is Privacy Worth?, 42 J. Legal. Stud. 249 (2013).

Monica Anderson & Andrew Perrin, 15% of Americans Don't Use the Internet: Who Are They? , Pew Research Center, July 28, 2015, available at <http://www.pewresearch.org/fact-tank/2015/07/28/15-of-americans-dont-use-the-internet-who-are-they/> .

Antilles Steamship Co. Ltd. v Members of American Hull Ins. Synd., 733 F2d 195 (2d Cir 1984).

Ian Ayres & Alan Schwartz, The No-Reading Problem in Consumer Contract Law, 66 Stan. L. Rev. 545 (2014).

Omri Ben-Shahar & Carl E. Schneider, More than You Wanted to Know: The Failure of Mandated Disclosure (Princeton 2014).

Omri Ben-Shahar & Lior Jacob Strahilevitz, A New Approach to Contract Interpretation: The Consumer Survey Method (unpublished manuscript 2016).

Russell Brandom, Someone's Trying to Gut America's Strongest Biometric Privacy Law, The Verge, May 27, 2016, available at <http://www.theverge.com/2016/5/27/11794512/facial-recognition-law-illinois-facebook-google-snapchat> .

Corley v. Google Inc., Complaint for Violations of the Electronic Communications Privacy Act, 18 U.S.C. § 2518 *et seq.*, Case No. 5-16-cv-00473, Jan. 27, 2016.

Shari S. Diamond & David J. Franklyn, Trademark Surveys: An Undulating Path, 92 Texas Law Rev. 2029-2073 (2014).

Wenyu Dou, Will Internet Users Pay for Online Content?, 44 J. Advertising Research 349 (2004).

Paul Ducklin, Automatic Photo Tagging: Facebook Friendships Get Creepier, Dec. 17, 2010 available at <https://nakedsecurity.sophos.com/2010/12/17/facebook-friendships-get-creepier/>.

Maeve Duggan, *The Demographics of Social Media Users*, Pew Research Center: Internet, Science & Tech., Aug. 19, 2015, available at <http://www.pewinternet.org/2015/08/19/the-demographics-of-social-media-users/> (reporting data from March and April of 2015).

Facebook Data Policy, Last Revised Jan. 30, 2015, Available at https://www.facebook.com/full_data_use_policy .

David A. Hoffman & Tess Wilkinson-Ryan, Legal Promise and Psychological Contract, 47 Wake Forest L. Rev. 843 (2012).

David A. Hoffman, From Promise to Form: How Contracting Online Changes Consumers, __ NYU L. Rev. __ (forthcoming 2017), unpublished draft available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2724661 (Jan. 29, 2016).

Brian Holland, Privacy Paradox 2.0, 19 Widener L.J. 893 (2010).

Jared Howe, How Much is Your Personal Data Worth?, Private WiFi, June 9, 2015, available at <http://blog.privatewifi.com/how-much-is-your-personal-data-worth/>

In re Facebook Biometric Information Privacy Litigation, Defendant Facebook, Inc.'s Motion to Dismiss, Case No. 3:15-cv-03747-JD, Oct. 9, 2015.

In re Google Inc. Gmail Litigation, Order Granting in Part and Denying in Part Defendant's Motion to Dismiss, Case No. 13-MD-02430-LHK, Sep. 26, 2013, available at 2013 WL 5423918.

In re Yahoo Mail Litigation, 7 F. Supp.3d 1016 (N.D. Cal. 2014).

Charles L. Knapp, Is There a "Duty to Read?," 66 Hastings L.J. 1083 (2015).

Matthew B. Kugler, Measuring Sponsorship Materiality, Dec. 27, 2015 Working Paper, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2628522.

Matthew B. Kugler & Lior Jacob Strahilevitz, Surveillance Duration Doesn't Affect Privacy Expectations: An Empirical Test of the Mosaic Theory, 2015 Sup. Ct. Rev. 205 (2016).

Matthew B. Kugler & Lior Jacob Strahilevitz, The Myth of Fourth Amendment Circularity (in progress, forthcoming 2017).

Licata v. Facebook, Inc., Consolidated Class Action Complaint, Case No. 3:15-cv-03747-JD, Aug. 28, 2015.

Stanislav Mamonov & Raquel Benbunan-Fich, An Empirical Investigation of Privacy Breach Perceptions Among Smartphone Application Users, 49 Computers in Hum. Behav. 427 (2015).

Florencia Marotta-Wurgler, Some Realities of Online Contracting, 19 Sup. Ct. Econ. Rev. 11 (2011).

Florencia Marotta-Wurgler, Does Contract Disclosure Matter, 168 J. Institutional & Theoretical Econ. 94 (2012).

Florencia Marotta-Wurgler, Understanding Privacy Policies: Content, Self-Regulation, and Markets, __ Journal of Legal Studies (forthcoming 2017) (this volume).

Kirsten Martin, Privacy Notices as Tabula Rasa: An Empirical Investigation into how Complying with a Privacy Notice is Related to Meeting Privacy Expectations Online, J. of Pub. Pol'y & Mktg. (forthcoming 2015).

Aleecia M. McDonald & Lorrie Faith Cranor, The Cost of Reading Privacy Policies, 4 I/S J.L. & Pol'y Info. Soc'y 543 (2008).

Ann Morales Olazabel, Howard Marmorstein and Dan Sarel, *Frequent Flyer Programs: Empirically Assessing Consumers' Reasonable Expectations*, 51 Am. Bus. L.J. 175 (2014).

Andrew Perrin & Maeve Duggan, Americans' Internet Access 2000-2015, Pew Research Center, June 26, 2015, available at <http://www.pewinternet.org/2015/06/26/americans-internet-access-2000-2015/>

Pezen v. Facebook, Class Action Complaint, Case No. 1:15-cv-03484, April 21, 2015.

Victoria C. Plaut & Robert P. Bartlett, III, Blind Consent? A Social Psychological Investigation of Non-Readership of Click-Through Agreements, 36 Law & Human Behav. 293 (2012).

Ariel Porat & Lior Jacob Strahilevitz, Personalizing Default Rules and Disclosure with Big Data, 112 Mich. L. Rev. 1417 (2014).

Lee Rainie, The State of Privacy in America: What We Learned, Pew Research Center, Jan. 20, 2016, available at <http://www.pewresearch.org/fact-tank/2016/01/20/the-state-of-privacy-in-america/>.

Joel R. Reidenberg et al., Disagreeable Privacy Policies: Mismatches Between Meaning and Users' Understanding, 2014 draft, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2418297.

Michael J. Seiler, Do Liquidated Damages Clauses Affect Strategic Mortgage Default Morality? A Test of the Disjunctive Thesis, Real Estate Econ. 1 (2016).

Jonathan Stempel, Google Won't Face Email Privacy Class Action, Reuters.com, Mar. 19, 2014, available at <http://www.reuters.com/article/us-google-gmail-lawsuit-idUSBREA2I13G20140319>.

Peter P. Swire, Financial Privacy and the Theory of High-Tech Government Surveillance, 77 Wash. U. L. Q. 461 (1999).

Yana Welinder, A Face Tells More than a Thousand Posts: Developing Face Recognition Privacy in Social Networks, 26 Harv. J. L. & Tech. 165 (2012).

Tess Wilkinson-Ryan, Legal Promise and Psychological Contract, 47 Wake Forest L. Rev. 843 (2012).

Tess Wilkinson-Ryan, A Psychological Account of Consent to Fine Print, 99 Iowa L. Rev. 1745 (2014).

Tess Wilkinson-Ryan & David A. Hoffman, The Common Sense of Contract Formation, 69 Stan. L. Rev. 1269 (2015).

Appendix – Text of Experimental Questions

I. Email:

A.) Prompts:

General instruction:

“Suppose that when you signed up with your current email provider you agreed that they could show advertisements next to your inbox in exchange for a free account. Suppose further than when signing up for the account, you read and agreed to the following terms and conditions:”

Then random assignment among these three prompts, all of which were taken from Gmail’s policy language:

- [Section 17] “Advertisements may be targeted to the content of information stored on the [email provider’s] services, queries made through the provider’s affiliated search engine, or other information”
- [Current language] “[Email provider’s] automated systems analyze your content (including emails) to provide you personally relevant product features, such as customized search results, tailored advertising, and spam and malware detection.”
- [Section 8] “[Email provider] reserves the right to pre-screen, review, flag, filter, modify, refuse or remove any or all content from any service. For some services, [email provider] may provide tools to filter out explicit sexual content.”

For the replication study, the options were Section 8, the current language, the below language from Yahoo:

- “[Email provider’s] automated systems analyze your content (including emails) to provide you personally relevant product features, such as customized search results, tailored advertising, and spam and malware detection.”

B.) Questions:

Would your agreement to this provision allow the email provider to direct its automated systems to scan the contents of the emails you send and receive and show you personalized advertisements? For example, if emails you send and receive regularly mention the words “tired” and “sleepy” the automated system might show you more ads from mattress sellers.

- They definitely would be allowed to do so. (1)
- They probably would be allowed to do so. (2)

- They probably would not be allowed to do so. (3)
- They definitely would not be allowed to do so. (4)

On a scale of 1 to 10, how intrusive is the email provider's automated email scanning and ad personalization practice?

If there were an option to keep the same email account but pay some amount of money to avoid having the automated systems analyze email content for the purposes of showing you personalized advertisements, how would you respond?

- I would keep the free email account with the automated email analysis and personalized advertisements.
- I would be willing to pay some amount of money to avoid the automated analysis.

[For those selecting 2 above] How much would you be willing to pay per year? ___ dollars ___ cents

Dollars (1)

Cents (2)

For the replication study,

- Half of the participants were asked "Why did you reach that answer?" after the "allow" question.
- And all were asked "How well do you feel you understood the passage?" on a 1 to 10 scale.
- Time on page was recorded.

II. Facebook:

A.) Prompts:

"The following language appears in the Data Policy on Facebook's web site:"

Random Assignment Among These Three Prompts:

1.)

"We collect the content and other information you provide when you use our Services, including when you sign up for an account, create or share, and message or communicate with others. This can include information in or about the content you provide, such as the location of a photo or the date a file was created. We also collect information about how you use our Services, such as the types of content you view or engage with or the frequency and duration of your activities.

We also collect content and information that other people provide when they use our Services, including information about you, such as when they share a photo of you, send a message to you, or upload, sync or import your contact information.

We are able to deliver our Services, personalize content, and make suggestions for you by using this information to understand how you use and interact with our Services and the people or things you're connected to and interested in on and off our Services.

We also use information we have to provide shortcuts and suggestions to you. For example, we are able to suggest that your friend tag you in a picture by comparing your friend's pictures to information we've put together from your profile pictures and the other photos in which you've been tagged. If this feature is enabled for you, you can control whether we suggest that another user tag you in a photo using the 'Timeline and Tagging' settings."

2.)

"We use the information we receive about you in connection with the services and features we provide to you and other users like your friends, the advertisers that purchase ads on the site, and the developers that build the games, applications, and websites you use. For example, we may use the information we receive about you:

- as part of our efforts to keep Facebook safe and secure;
- to provide you with location features and services, like telling you and your friends when something is going on nearby;
- to measure or understand the effectiveness of ads you and others see;
- to make suggestions to you and other users on Facebook, such as: suggesting that your friend use our contact importer because you found friends using it, suggesting that another user add you as a friend because the user imported the same email address as you did, or suggesting that your friend tag you in a picture they have uploaded with you in it.

Granting us this permission not only allows us to provide Facebook as it exists today, but it also allows us to provide you with innovative features and services we develop in the future that use the information we receive about you in new ways .We are able to suggest that your friend tag you in a picture by comparing your friend's pictures to information we've put together from the photos you've been tagged in. You can control whether we suggest that another user tag you in a photo using the 'How Tags work' settings."

3.)

“When someone uploads a photo of you, we might suggest that they tag you in it. We’re able to compare your friend’s photos to information we’ve put together from your profile pictures and the other photos you’re tagged in. If this feature is turned on for you, you can choose whether or not we suggest your name when people upload photos of you. Adjust this in your Timeline and Tagging settings.

We currently use facial recognition software that uses an algorithm to calculate a unique number (‘template’) based on someone’s facial features, like the distance between the eyes, nose and ears. This template is based on your profile pictures and photos you’ve been tagged in on Facebook.

We use these templates to help you tag photos by suggesting tags of your friends. If you remove a tag from a photo, that photo is not used to create the template for the person whose tag was removed. We also couldn’t use a template to recreate an image of you.”

For the replication study, the options were the first and third of the above, and a modification of the third, “when someone uploads” that ended:

We use these templates to help you tag photos by suggesting tags of your friends. If you remove a tag from a photo, that photo is not used to create the template for the person whose tag was removed. We automatically delete all facial recognition information once it has been stored in our system for three years. We also couldn’t use a template to recreate an image of you.”

B.) Questions

Suppose you had previously read the Data Policy’s / Help Center’s language and had not adjusted your Timeline and Tagging settings. Suppose further that the following scenario occurs: A Facebook friend of yours uploads a photo of you and them to Facebook. Because Facebook already has analyzed other photos of you, its facial recognition software suggests to your friend that you be tagged (captioned) in the photo, and your friend agrees to tag you in the photo.

	Yes (1)	No (2)
Did Facebook's language (above) inform you that information about your facial features was being collected and stored?	<input type="radio"/>	<input type="radio"/>
Did Facebook's language (above) inform you of the reason why information about your facial features was being collected, stored, and used?	<input type="radio"/>	<input type="radio"/>
Did Facebook's language (above) inform you of the length of time for which information about your facial features would be stored?	<input type="radio"/>	<input type="radio"/>
Would your decision not to adjust your Timeline and Tagging settings allow Facebook to collect, store, and use information about your facial features?	<input type="radio"/>	<input type="radio"/>

On a scale of 1 to 10, how intrusive is Facebook's use of facial recognition software to suggest tags for people whose faces appear in uploaded photos?

For the replication study,

- All were asked "How well do you feel you understood the passage?" on a 1 to 10 scale.
- Time on page was recorded.