A Field Experiment on Biased Beliefs and Information Overload in Consumer Privacy

Ignacio N. Cofone* and Adriana Z. Robertson**

Abstract

While many scholars call for companies to implement consumer privacy notices to increase transparency, others suggest that notices are ineffective at increasing consumer awareness of how their personal information is managed. We find whether the reason for the ineffectiveness of privacy notices is a common behavioral bias: the non-belief in the law of large numbers (NBLLN). People affected by NBLLN misunderstand how quickly an observer can piece together clues based on available information. There is strong evidence that NBLLN is among the most prevalent behavioral factors in the population. This work builds on our existing research project about measuring privacy harms and on consumer valuations of privacy. Using a field experiment where consumers receive different types of privacy notifications, we will test the effectiveness of privacy notifications designed to account for the fact that consumers have difficulty correctly aggregating large amounts of information.

* Research Fellow, NYU Information Law Institute. ignacio.cofone@nyu.edu.

** Assistant Professor, University of Toronto Faculty of Law and Rotman School of Management. adriana.robertson@utoronto.ca.

Extended Abstract

There is an important contradiction in the literature on the effectiveness of consumer privacy notices. While many scholars call for more companies to implement consumer privacy notices as a way to increase transparency,[1] others suggest that notices are ineffective at increasing consumer awareness of how their personal information is managed.[2] Indeed, empirical evidence has shown that simplifying disclosures has no effect on consumer awareness, suggesting that complexity in language is not the main driver.[3] Moreover, other empirical work suggests that the language used in a privacy policy is irrelevant, which in turn suggests that consumers do not react to different kinds of language.[4]

We propose that an explanation for the ineffectiveness of privacy notices is a common behavioral bias called the "non-belief in the law of large numbers" (NBLLN). In essence, people affected by NBLLN misunderstand how quickly an observer can piece together clues based on available information.

---

1 Ryan Calo, *Against Notice Skepticism in Privacy (and Elsewhere)*, 87 NOTRE DAME L.REV. 1027 (2011) (proposing visceral notices for privacy); Paula J. Dalley, *The Use and Misuse of Disclosure as a Regulatory System*, 34 FLA. STATE UNIV. L.REV. 1089 (2006) (noting the provision of notices as a common method for regulation); William M. Sage, *Regulating through Information: Disclosure Laws and American Health Care*, 99 COLUMBIA L.REV. 1701–1829 (1999) (explaining the provision of notices as a common method for regulation in medicine).

2 Kirsten Martin, *Do Privacy Notices Matter? Comparing the Impact of Violating Formal Privacy Notices and Informal Privacy Norms on Consumer Trust Online*, 45 J. LEG. STUD. 191 (2016) (using a vignette study to show that formal privacy notices actually reduce consumer trust on a website). See also Solon Barocas & Helen Nissenbaum, *On Notice : The Trouble with Notice and Consent* (2009); Aleecia McDonald & Lorrie Faith Cranor, *The Cost of Reading Privacy Policies*, 4 J. L.POLICY INF. SOC. 543 (2008) (showing the time and energy needed to comprehend privacy policies); Susanna Kim Ripken, *The Dangers and Drawbacks of the Disclosure Antidote: Toward a More Substantive Approach to Securities Regulation*, 58 BAYL. L.REV. 139 (2006) (explaining the limits of a disclosure-based policy generally and suggesting direct conduct regulation through the example of securities).

3 Omri Ben-Shahar & Adam Chilton, *Simplification of Privacy Disclosures: An Experimental Test*, 45 J. LEG. STUD. 41 (2016) (finding that best-practice simplification techniques have little or no effect on respondents' comprehension of disclosures).

4 Lior Jacob Strahilevitz & Matthew B. Kugler, *Is Privacy Policy Language Irrelevant to Consumers?*, 45 J. LEG. STUD. 69 (2016) (testing language in privacy policies).

There is strong evidence that NBLLN is among the most prevalent behavioral factors in the population. For example, in a recent large-scale study involving 1500 participants drawn from a representative panel of US adults, Stango et al. studied the incidence of 17 different widely recognized "behavioral factors," including NBLLN.[5] They report that fully 87% of the participants in their study exhibited NBLLN, the most of any factor in the study.[6] This suggests that a huge proportion of the population is vulnerable to this bias. When faced with large amounts of data, the average person is bad at estimating the informativeness of each piece of newly arriving information.[7] In the context of one's online data, this means that individuals will tend to underestimate the amount of privacy that they are ceding to commercial parties in each interaction. The upshot of this is that individuals will tend to give away their private data too easily, or sell it too cheaply.

Prior literature has advanced our understanding of consumer behavior by identifying the prevalence of behavioral biases in consumer's privacy choices.[8] While the effect of NBLLN on individual willingness to sell private data is not limited to the internet context, the problem is particularly acute in the digital domain. The sheer amount of data that can feasibly be collected in the digital world is exponentially larger than in the analogue world. Individuals afflicted by NBLLN will therefore have a particularly hard time understanding just how valuable their digital data is. Instead, they suffer from information overload. As storage and processing costs continue to fall, and as machine learning algorithms become more effective, the amount of information that can be extracted from each data point will only increase, exacerbating this problem further.[9]

---

5 VICTOR STANGO, JOANNE YOONG & JONATHAN ZINMAN, THE QUEST FOR PARSIMONY IN BEHAVIORAL ECONOMICS: NEW METHODS AND EVIDENCE ON THREE FRONTS (2017), http://www.nber.org/papers/w23057.

6 *Id.* at 61. This makes NBLLN the second most prevalent behavioral factor from among these 17 factors, after violation of the general axiom of revealed preference ("GARP").

7 Daniel Kahneman & Amos Tversky, *Subjective probability: A judgment of representativeness*, 3 COGNIT. PSYCHOL. 430, 444 (1972) (famously observing that "[t]he notion that sampling standard deviation decreases in proportion to sample size is apparently not part of man's repertoire of intuitions.").

8 Acquisti, Brandimarte, and Loewenstein, *Privacy and Human Behavior in the Age of Information*, 347 SCIENCE 509 (2015).

9 *See infra* notes 9-11 and accompanying text for a brief discussion of the formal model of this phenomenon.

An example can help to clarify this phenomenon. Suppose that an individual named Abby is faced with a decision about whether to use a new smartphone application. She knows that, as a condition of using it, she will be granting to a company, Poodle, the creator of the application, access to the data she produces while using it. This grant of her private data is the price she pays for the application – she effectively sells her data in return for access to the application. While she realizes that Poodle can use this information to learn about her, if she suffers from NBLLN, she will underestimate *how much* Poodle can learn about her. In particular, she will underestimate the significance of such data when combined with the data that she shares, through a different application, with Goggles. In other words, she will underestimate her privacy loss, causing her to undervalue her private data. Moreover, this effect will be significantly enhanced if Poodle and Goggles combine their information about her. As a result, she will be willing to sell her data too cheaply, and may therefore mistakenly agree to grant Poodle access to her data.

In prior work, we proposed a model to formalize the concept of privacy loss based on Bayesian updating.[10] We then used the cognitive bias to build on this formal model by formalizing the application of Benjamin et al.'s model of NBLLN[11] to our model of privacy loss.[12] Adding this element causes the agent to undervalue her personal data. As a result, she will sell too much of her data at too low of a price. Based on our theoretical framework, we then made concrete policy proposals of what an effective privacy notification should look like.[13] In this paper, we will empirically test those proposals. With this test, we will be able to provide a unified and testable framework for interpreting the existing empirical results on the effectiveness of notices. To our knowledge, this will be the first empirical study of the issue of information aggregation by consumers, as well as the first attempt to test methods of de-biasing consumers with NBLLN.

---

[10] Ignacio Cofone & Adriana Z. Robertson, *Privacy Harms*, 69 HASTINGS L.J. ___ (forthcoming 2018).

[11] Daniel J. Benjamin, Matthew Rabin & Collin Raymond, *A Model of Nonbelief in the Law of Large Numbers*, 14 J. EUR. ECON. ASSOC. 515 (2016).

[12] Ignacio Cofone & Adriana Z. Robertson, *Consumer Privacy in a Behavioral World*, 69 HASTINGS L.J. ___ (forthcoming 2018).

[13] *Id.* These effective privacy notifications can either be mandated through disclosure requirements (for example, by the FTC) or could simply become "best practices" in the industry.

All participants will be told that they are completing a survey about consumer behavior, and will be presented with the same scenario involving a new smartphone application.[14] Participants will then be randomly assigned into one of three groups. The three groups will receive different privacy disclosure statements. Group 1 will receive a "standard" disclosure statement, which is consistent with the types of disclosure statements that exist in the market. Group 2 will receive a "simplified" disclosure statement, consistent with the best-practice simplified disclosure statements examined by Ben-Shahar and Chilton.[15] Group 3 will receive an "NBLLN-compliant" disclosure statement, based on the principles outlined in prior work by Cofone and Robertson.[16]

The crucial part of the experiment is the drafting of both the scenario and the three disclosure statements. As such, we expect to pilot several versions of the language and to solicit feedback from a variety of colleagues before settling on the final wording.[17] We will pilot the wording using Amazon's "Mechanical Turk" (or "MTurk") platform, which has been successfully used for experiments in law[18] as well as other social sciences,[19] and which provides a fast and low-cost way to test our draft wording. During the pilot phase, we will look for evidence that the respondents understand the *wording* of the scenario and the disclosure statement, as well as evidence of what they believe the survey is asking about. We will

---

14 Our intention is that the description be a correct description of what the respondents will be doing, while at the same time be sufficiently vague that the respondents are not alerted to the fact that the real object of interest is their response to the privacy disclosure, and not to other aspects of the smartphone application. This will enable us to avoid using outright deception, while at the same time allowing us to collect unpolluted responses.

15 Ben-Shahar & Chilton, *supra* note 3.

16 Cofone & Robertson, *supra* note 12.

17 *See* James J. Choi & Adriana Z. Robertson, *What Matters to Individual Investors? Evidence from the Horse's Mouth*, working paper (2017) (discussing the repeated process of drafting, piloting, and soliciting feedback on wording before settling on the final text of the questions for a large-scale survey).

18 *See, e.g.*, Jane Bambauer, Jonathan Loe and D. Alex Winkelman, *A Bad Education*, 2017 U. ILL. L. REV. 109 (2017) (conducting an experiment using participants recruited with MTurk). *See also* Emily Satterthwaite, *Can Audits Encourage Tax Evasion?" An Experimental Assessment*, 20 FL. TAX REV. 1 (2016) (conducting an experiment on tax evasion using MTurk).

19 *See, e.g.,* Augustin Landier, Yueran Ma and David Thesmar, *New Experimental Evidence on Expectations Formation*, working paper (2017) (using participants recruited from MTurk to experimentally test theories of expectation formation).

not be looking for evidence about whether they understand the *implications* of the disclosure statements, as this will be the purpose of the experiment. While we may also use MTurk for the experiment, we continue to explore other alternative venues (such as Qualtrics) to determine which one fits the design the best. We are particularly interested in the workshop participants' opinion on this point.

After having read the scenario and the appropriate disclosure, participants to the experiment will be asked how likely they are to use the smartphone application. We hypothesize that respondents in Group 3 will be less likely to indicate that they would use the smartphone app than respondents in Groups 1 or 2. This is consistent with the joint hypothesis that individuals care about their privacy, but do not understand the privacy implications of their actions because of NBLLN.

We will also ask the participants questions designed to test their understanding of the implications of the disclosure. These questions will take the following form: "based on the disclosure you just read, which of the following would the company be permitted to do?" We hypothesize that respondents in Group 3 will be more likely to correctly identify which activities are permitted and which are prohibited under the terms of the disclosure than respondents in Groups 1 or 2. This represents a direct test of the hypothesis that NBLLN compliant disclosures can assist consumers in comprehending privacy implications.

Finally, we will ask a third set of questions to the participants who indicate that they would have accepted the privacy disclosure. These questions will take the following form: "if you knew that the maker of the smartphone application was allowed to [X], would you still have consented to the privacy policy?" In many cases, X will consist of something that *would actually be* permitted under the privacy agreement. The degree to which respondents answer inconsistently – that is, they indicate that they *would not* consent to X, while they had *in fact* consented to X – provides a direct way to measure the respondent's privacy error. We hypothesize that the rate of errors will be lower for respondents in Group 3 than it is for those in Groups 1 and 2.