

Articles

Consumer Privacy in a Behavioral World*

IGNACIO N. COFONE** & ADRIANA Z. ROBERTSON***

On March 28, 2017, Congress killed the FCC's attempt to protect consumer privacy on the internet and allowed ISPs to continue to track their users' online behavior. We evaluate the impact of this decision for consumer privacy in light of biased beliefs and information overload. We do so through a well-documented behavioral bias: Non-belief in the Law of Large Numbers. In doing so, we provide a framework for protecting consumer privacy. We then suggest private law and regulatory solutions to do so in a more effective way than either the current or the now-repealed regime.

* We would like to thank Ian Ayres, Jack Balkin, Rebecca Crootof, Joshua Fairfield, Claudia Haupt, Edward Iacobucci, Christine Jolls, Ido Kilovaty, M. Henry Linder, Mason Marks, Katherine Strandburg, Eduardo Stordeur, and Ari Waldman for their helpful feedback, as well as two referees and the editorial staff at the *Hastings Law Journal*. This paper also benefited from comments received at the Canadian Law & Economics Association annual conference (Toronto, 2017) and an internal presentation at the Yale Law School Information Society Project.

** Post-doctoral 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.

TABLE OF CONTENTS

INTRODUCTION.....	1472
I. CONSUMER PRIVACY IN THE DIGITAL WORLD	1476
A. THE FCC PRIVACY ORDER	1476
B. PRIVACY IMPLICATIONS OF ISPs	1481
C. THE FCC PRIVACY ORDER AS A DEFAULT RULE	1484
D. ECONOMIC ARGUMENTS FOR INTERVENTION.....	1486
E. THE PRIVACY PARADOX AND NBLLN.....	1487
II. PRIVACY LOSS IN THE FACE OF NBLLN.....	1488
A. OUR MODEL OF PRIVACY LOSS.....	1488
1. <i>Types and Clues</i>	1488
2. <i>Aggregating Clues</i>	1489
B. NON-BELIEF IN THE LAW OF LARGE NUMBERS.....	1491
1. <i>A Well-Known Bias</i>	1491
2. <i>Addressing Concerns About Behavioral Economics</i> ..	1492
C. NBLLN AND PRIVACY LOSS	1493
1. <i>Formalization</i>	1493
2. <i>Intuition</i>	1496
III. ADDRESSING BIASED BELIEFS AND INFORMATION OVERLOAD	1497
A. IMPLICATIONS OF THE MODEL	1497
B. PRIVATE LAW APPROACHES	1498
C. REGULATORY APPROACHES	1501
1. <i>Starting Point: Clear and Simple Disclosures</i>	1503
2. <i>How Information Can Be Combined</i>	1503
3. <i>Classes of Information</i>	1505
4. <i>Impact on Existing Notices Literature</i>	1506
CONCLUSION	1507

INTRODUCTION

Advertising is the lifeblood of the internet. Two of the largest players in the online world—Alphabet and Facebook—jointly earned 151.6 billion dollars in advertising revenues in 2017. Alphabet, for its part, is sometimes characterized as an advertising company rather than as a search company. A major difference between digital advertising and its more traditional cousin is the ability to personalize content. As an old adage in marketing goes, “half of advertising spending is wasted. The problem is that nobody knows which half.”¹ A company with more

1. See George Bradt, *Wanamaker Was Wrong—The Vast Majority of Advertising is Wasted*, FORBES (Sept. 14, 2016, 6:56 AM), <https://www.forbes.com/sites/georgebradt/2016/09/14/wanamaker-was-wrong-the-vast-majority-of-advertising-is-wasted/#28146de6483b>

information about consumers' preferences can better target its advertising, offering more products and services that are tailored to each consumer's preferences and interests. This ability to tailor produces more sales and higher payments for advertising spots compared to non-targeted advertising, increasing revenue for sellers. In turn, it allows sellers of advertising space—such as Facebook and Alphabet—to charge a premium.

For targeted advertising to work, companies need to track people's online behavior. One well-known way to do this is through the use of http cookies, which are small pieces of text sent back and forth between a server and a user's internet browser.² While cookies may be the best-known form of online tracking, websites have other means at their disposal, such as fingerprinting (in which a link contains embedded data showing how a user accessed a particular website)³ or by monitoring visitors' IP addresses. The ability to track online user behavior is not unique to websites. Internet Service Providers (ISP), in particular, can monitor everything their clients do online while connected to the ISP's server.⁴ ISP tracking begins "at the source" and, therefore, has the potential to be far more comprehensive than the tracking methods used by individual websites.⁵

However, on March 28, 2017, Congress disapproved the FCC's latest attempt to protect consumer data privacy on the Internet.⁶ The Obama

(John Wanamaker supposedly said "Half the money I spend on advertising is wasted; the trouble is I don't know which half.").

2. *HTTP Cookie*, MCGILL SCH. COMPUTER SCI., http://cs.mcgill.ca/~rwest/wikispeedia/wpcd/wp/h/HTTP_cookie.htm (last visited July 29, 2018).

3. Jonathan R. Mayer & John C. Mitchell, *Third-Party Web Tracking: Policy and Technology*, 2012 IEEE SYMP. ON SECURITY & PRIVACY 413, 420 (2012), <https://ieeexplore.ieee.org/stamp/stamp.jsp?tp=&arnumber=6234427>.

4. See Paul Ohm, *The Rise and Fall of Invasive ISP Surveillance*, 2009 U. ILL. L. REV. 1417, 1420 (2009) ("Everything we say, hear, read, or do on the Internet first passes through ISP computers. If ISPs wanted, they could store it all, compiling a perfect transcript of our online lives. In fact, nothing in society poses as grave a threat to privacy as the ISP, not even Google, a company whose privacy practices have received an inordinate amount of criticism and commentary. Although Google collects a vast amount of personal information about its users, an ISP can always access even more because it owns and operates a privileged network bottleneck, the only point on the network that sits between a user and the rest of the Internet. Because of this fact about network design, a user cannot say anything to Google without saying it first to his ISP, and an ISP can also hear everything a user says to any other websites like Facebook or eBay, things said that are unobtainable to Google. The potential threat to privacy from unchecked ISP surveillance surpasses every other threat online.").

5. See *id.* at 1438 ("In modern connected life, almost no other entity can access as much personal information . . . Because the ISP is the gateway—the first hop—to the Internet, almost any communication sent to anybody online is accessible first by the ISP . . . In fact, no other online entity can watch every one of a user's activities, making the ISP's viewpoint uniquely broad.").

6. See S.J. Res. 34, 115th Cong. (2017) (a joint resolution "providing for congressional disapproval under chapter 8 of title 5, United States Code, of the rule submitted by the Federal Communications Commission relating to 'Protecting the Privacy of Customers of Broadband and Other Telecommunications Services.'"). The resolution was passed by the House in a 215–205 vote,

Administration had adopted rules that limited this type of tracking. In essence, those rules would have required ISPs to obtain explicit opt-in consent to access, use, or sell certain types of users' personal information.⁷ Congress disapproved the new rules before they came into force.⁸

In recent years, there has been a growing chorus calling for further regulation of consumer data.⁹ Concerns about consumer privacy are so prominent that some scholars have argued that the entirety of information privacy law has been subsumed by consumer contract law, pushing aside other issues such as privacy torts and the Fourth Amendment.¹⁰ The FCC's order, in particular, received more than a quarter of a million filings, almost all of which supported adopting stronger consumer privacy rules.¹¹

A presumption in classical economic theory is that, in the absence of transaction costs or market failures, free exchanges between rational self-interested parties are mutually beneficial and will lead to an efficient allocation of resources. In theory, this argument should apply just as well to consumer data as it does to anything else.¹² It would follow, then, that

after being passed by the Senate in a 50–48 vote, both along party lines. *Actions Overview S.J. Res. 34–115th Congress (2017–2018)*, CONGRESS.GOV, <https://www.congress.gov/bill/115th-congress/senate-joint-resolution/34/actions> (last visited July 29, 2018). See generally Cecilia Kang, *Broadband Providers Will Need Permission to Collect Private Data*, N.Y. TIMES (Oct. 27, 2016), <https://www.nytimes.com/2016/10/28/technology/fcc-tightens-privacy-rules-for-broadband-providers.html> (explaining the FCC proposed rule for a general audience). Note that sometimes, even if broadband providers are listed as the subject of the regulation, the resolution specifically covered all ISPs, including mobile and other internet providers. See 81 Fed. Reg. 87,274, 87,334 (Dec. 2, 2016).

7. In particular, the rule would have required ISPs to obtain consumer consent before tracking financial, health, and children's data, and browsing history specifically for behavioral advertising purposes. See David Shepardson, *Trump Signs Repeal of U.S. Broadband Privacy Rules*, REUTERS (Apr. 3, 2017, 4:50 PM), <https://www.reuters.com/article/us-usa-internet-trump-idUSKBN1752PR> (reporting on the repeal of the FCC Order).

8. See S.J. Res. 34, 115th Cong. (2017). In that resolution, Congress disapproved of the FCC's rule relating to "Protecting the Privacy of Customers of Broadband and Other Telecommunications Services." *Id.* The FCC's rule included a timeline for stipulating when the different obligations would come into force, spanning from "90 days" to "twelve months," with an additional "twelve-month" period for small providers. 81 Fed. Reg. 87,274, 87,319, 87,341, 87,342 (Dec. 2, 2016).

9. See Brian X. Chen, *What the Repeal of Online Privacy Protections Means for You*, N.Y. TIMES (Mar. 29, 2017), <https://www.nytimes.com/2017/03/29/technology/personaltech/what-the-repeal-of-online-privacy-protections-means-for-you.html> (reporting on the repeal's significance for consumer protection and explaining these demands regarding the repeal of the FCC rules).

10. See Omri Ben-Shahar & Lior Jacob Strahilevitz, *Contracting over Privacy: Introduction*, 45 J. LEGAL STUD. S1, S1 (2016) (highlighting "a quiet legal transformation whereby the entire area of data privacy law has been subsumed by consumer contract law").

11. Protecting the Privacy of Customers of Broadband and Other Telecommunications Services, 31 FCC Rcd 2,500, 2,635 (Clyburn, Comm'r, approving in part and concurring in part).

12. Indeed, consumer data has been described as being "to this century what oil was to the last one." *Data Is Giving Rise to a New Economy*, ECONOMIST (May 6, 2017), <http://www.economist.com/news/briefing/21721634-how-it-shaping-up-data-giving-rise-new-economy> (arguing that data is a crucial fuel of the modern economy).

any economic argument for the regulation of the market for consumer data must begin by showing that (at least) one of these elements fails to hold. We take up this challenge.

We argue that the existence of a well-established behavioral bias justifies regulation of consumer data.¹³ To do so, we augment the formal model developed in previous work¹⁴ with a cognitive bias known as the “non-belief in the law of large numbers” (NBLLN). We are not the first to note that consumers may be unable to accurately estimate the marginal effects that their decisions have on their level of privacy. For example, Strandburg has argued that uncertainty might prevent individuals from accurately perceiving their level of privacy loss.¹⁵ While we agree with this general conclusion, our analysis focuses on the effect of a specific, well-defined, behavioral bias.

People affected by this bias suffer from a form of information overload, causing them to misunderstand how quickly an observer can piece together clues based on available information. Adding this type of information overload to our model of an otherwise rational and self-interested agent causes the agent to undervalue her personal data. As a result, she will sell too much of her data at too low a price, leading to an inefficiently low level of consumer privacy in the economy. This result follows without assuming any other transaction costs or other market failures.

In addition to providing a rationale for the regulation of consumer data, our proposal provides regulators with insight into how best to regulate this market. In particular, it suggests that, counterintuitively, changing defaults from an “opt-out” to an “opt-in” model is unlikely to have much of an effect. On the other hand, mandatory disclosures

13. This is not to say that there cannot also be rational reasons why individuals might give up more data than they would in a very simple market with no externalities. For example, suppose that a customer believes that the online party—such as her ISP—can use data from its other consumers to make highly reliable inferences about her. In this case, the incremental loss of her privacy from giving up her own data might be quite small. In this example, because the online party can make strong inferences across individuals, there is a negative externality from each individual’s decision to give up data. This negative externality leads to an inefficiently high level of data sharing. We thank Ed Iacobucci for this insight.

14. Ignacio N. Cofone & Adriana Z. Robertson, *Privacy Harms*, 69 HASTINGS L.J. 1039, 1049–58 (2018).

15. See Katherine J. Strandburg, *Free Fall: The Online Market’s Consumer Preference Disconnect*, 2013 U. CHI. LEGAL F. 95, 96 (2013) (“Internet users do not know the ‘prices’ they are paying for products and services supported by behavioral advertising because they cannot reasonably estimate the marginal disutility that particular instances of data collection impose on them.”); see also Jeff Sovern, *Opting In, Opting Out, or No Options At All: The Fight for Control of Personal Information*, 74 WASH. L. REV. 1033, 1072–74 (1999) (addressing the problems that the behavioral advertising business model involves for internet consumers seeking to properly estimate the costs of data collection).

regarding how informative particular pieces of data will be can help to increase the overall efficiency of the market for consumer data.

Our approach also helps to address a related concept known as the privacy paradox. This paradox can be summarized by the following question: why is it that individuals consistently indicate that they value privacy, while simultaneously giving their privacy away for almost nothing?¹⁶ We show that a consumer affected by NBLLN will act in a manner that is consistent with the privacy paradox.

The remainder of our argument proceeds as follows. In Part I, we survey the landscape of consumer privacy in the digital world. We begin by discussing prior policy interventions in the internet privacy space, including the recently disapproved FCC internet privacy rules. We also discuss some of the major economic arguments for and against regulatory interventions. Finally, we introduce the well-known privacy paradox, and argue that NBLLN can help to reconcile this paradox.

In Part II, we present a model of privacy loss and introduce our formal model of NBLLN. We then show how NBLLN can lead individuals to undervalue their personal data. This undervaluation, and the mispricing that follows from it, is particularly acute in the context of tracking by ISPs, and provides an economic argument for government intervention in the digital privacy space.

In Part III, we discuss the implications of these findings for policy makers. After surveying potential private law solutions, we argue that regulatory solutions are more appropriate in overcoming NBLLN. We then suggest how to shape future regulatory solutions in the face of NBLLN and offer specific policy recommendations.

I. CONSUMER PRIVACY IN THE DIGITAL WORLD

A. THE FCC PRIVACY ORDER

On December 2, 2016, the FCC took an unprecedented step towards protecting consumer data from the prying eyes of their Internet Service Providers. It did so by publishing an Order, entitled “Protecting the Privacy of Customers of Broadband and Other Telecommunications Services,” (“FCC Privacy Order”) extending “traditional [regulatory]

16. See Patricia A. Norberg et al., *The Privacy Paradox: Personal Information Disclosure Intentions Versus Behaviors*, 41 J. CONSUMER AFF. 100, 100 (2007) (exploring “the privacy paradox,” which refers to “the relationship between individuals’ intentions to disclose personal information and their actual personal information disclosure behaviors.”); see also Alessandro Acquisti et al., *The Economics of Privacy*, 52 J. ECON. LITERATURE 1, 39–41 (2016) (surveying the economics of privacy and reviewing the literature on the privacy paradox).

privacy requirements” to ISPs.¹⁷ The FCC had previously classified broadband internet access as telecommunications in its Open Internet Order.¹⁸ According to the FCC, this meant that ISPs were subject to section 222 of the Communications Act, which requires telecommunications carriers to protect their users’ personal information.¹⁹ The FCC Privacy Order then applied the consumer privacy requirements of the Communications Act to ISPs.²⁰ In doing so, the FCC Privacy Order added ISPs to the group of companies obligated to abide by a set of duties. We classify these duties into two categories: “transparency duties” and “consent duties.”

The transparency duties centered largely around disclosure to consumers and regulators. For example, ISPs would have been required to incorporate breach notifications, as well as persistent notices about the information collected, including information on how the data could be used and with whom it could be shared.²¹ This included notifying consumers about the types of information being collected, how the information was being collected, the purposes for which the ISP would use or share the information, and the types of entities with whom the ISP would share the information. ISPs would also have been required to take reasonable measures to secure consumers’ personal information based on guidelines set out by the FCC.²² The transparency rules also prohibited ISPs from requiring consumers to waive their privacy rights as a condition of service. This would have prevented ISPs from offering a

17. Protecting the Privacy of Customers of Broadband and Other Telecommunications Services, 81 Fed. Reg. 87,274 (Dec. 2, 2016) (now codified at 47 C.F.R. pt. 64) [hereinafter *FCC Privacy Order*].

18. See Preserving the Open Internet, 76 Fed. Reg. 59,192, 59,214 (Sept. 23, 2011) (now codified at 47 C.F.R. pt. 8) (establishing non-discrimination online through § 8.1 on purpose, § 8.3 on transparency, § 8.5 on blocking, and § 8.7 on no unreasonable discrimination); see also Appropriate Framework for Broadband Access to the Internet over Wireline Facilities, 20 FCC Rcd 14,986, 14,988 (Dortch, Sec’y, Policy Statement). (establishing the four principles of the open internet: (i) “consumers [deserve] access to the lawful Internet content of their choice”; (ii) “consumers [should be allowed] to run applications and use services of their choice, subject to the needs of law enforcement”; (iii) “consumers [should be able] to connect their choice of legal devices that do not harm the network”; and (iv) “consumers [deserve to choose their] network providers, application and service providers, and content providers of choice”).

19. Communications Act of 1934, 47 U.S.C. § 222(a) (2008).

20. *FCC Privacy Order*, *supra* note 17, at 87,328–87,329, 87,333, 87,343–87,344. Note that the rule did not include services that the FTC has authority over (like websites, mobile applications, and other services of broadband providers), government surveillance, or law enforcement activities. See generally *FCC Privacy Order*, *supra* note 17, 87,274, 87,276 (“In this Report and Order (Order) we apply the privacy requirements of the Communications Act of 1934, as amended (the Act) to the most significant communications technology of today—broadband Internet access service (BIAS).”).

21. *FCC Privacy Order*, *supra* note 17, at 87,312–87,315 (indicating that affected consumers and the FCC had to be notified, and, under certain circumstances, the FBI and the Secret Service should have been notified as well).

22. *FCC Privacy Order*, *supra* note 17, at 87,327–87,328 (providing guidelines (rather than a checklist) on how to operationalize this and relied on FTC best practices).

“take it or leave it” contract, which would have left consumers with no choice but to consent to their data being used or shared for commercial purposes to receive service.

In contrast, the consent duties were focused on obtaining either opt-in or opt-out consent from consumers in various contexts. ISPs would have been required to obtain opt-in consent in order to use or share sensitive information, and to obtain opt-out consent in order to use or share non-sensitive information. The rules established wide categories for sensitive information, such as geolocation, financial information, health information, children’s information, social security numbers, web browsing history, app usage history and content of communications. Non-sensitive information, in contrast, was defined as all personally identifiable information that did not fall under one of the categories captured by the definition of sensitive information, such as consumers’ contact information. The consent duties also required ISPs to provide enhanced notice and obtain affirmative consent in order to use consumers’ personal information in exchange for financial incentives.²³

According to the FCC’s news release, the rules aimed to “ensure broadband customers have meaningful choice, greater transparency and strong security protections for their personal information collected by ISPs.”²⁴ The idea was that, by broadening the framework of consumer consent to ISPs and calibrating it to information sensitivity, the FCC Privacy Order would regulate them in a manner consistent with the FTC privacy rules and the Administration’s Consumer Privacy Bill of Rights.²⁵

The consent rules were both more important and more controversial than the transparency rules.²⁶ By the end of January 2017—before

23. *FCC Privacy Order*, *supra* note 17, at 87,317–87,318. This provision included heightened disclosure duties for the “pay for privacy” plans, and also required explicit affirmative consent from consumers. The FCC would have analyzed these programs on a case-by-case basis.

24. Press Release, Fed. Comm’n Comm’n, FCC Adopts Privacy Rules to Give Broadband Consumers Increased Choice, Transparency and Security for Their Personal Data (Oct. 27, 2016), https://apps.fcc.gov/edocs_public/attachmatch/DOC-341937A1.pdf.

25. *Id.*

26. *See, e.g.*, Protecting the Privacy of Customers of Broadband and Other Telecommunications Services, 31 FCC Rcd 2,500, 2,638 (Pai, Comm’r, dissenting) (agreeing with his fellow commissioners’ previous statements that “consumers should not have to be network engineers to understand who is collecting their data and they should not have to be lawyers to determine if their information is protected,” but finding that the FCC’s order dramatically departed from this principle); *see also* THE FED. COMM’N COMM’N, WC DOCKET NO. 16-106, FACT SHEET: THE FCC ADOPTS ORDER TO GIVE BROADBAND CONSUMERS INCREASED CHOICE OVER THEIR PERSONAL INFORMATION (2016), <https://www.fcc.gov/document/fcc-releases-rules-protect-broadband-consumer-privacy> (last visited July 29, 2018). There were some exceptions to the duty to obtain consent before using personal information, although they were limited. *Id.* at 2. Only non-sensitive information (information that would necessarily be provided to the broadband service), information used to provide and market services and equipment (such as use of a modem) typically marketed with the broadband, and information needed to protect the ISP from fraudulent use of its network were

Congress disapproved the rule—the FCC had received almost a quarter of a million filings from individuals in support of the FCC Privacy Order, as well as eleven petitions to reconsider.²⁷ Nine associations submitted a joint petition for a stay,²⁸ which was opposed by a group of eleven consumer associations in February 2017.²⁹

ISPs argued that the FCC Privacy Order would be both costly and burdensome to implement.³⁰ More importantly, the ISPs maintained that they had already devised a set of voluntary “privacy and data security principles,” which “include a commitment to take reasonable measures to protect customer information from unauthorized use, disclosure, or access, taking into account the nature and scope of their activities, the sensitivity of the data, the size of the ISP, and technical feasibility.”³¹

Acting pursuant to the Congressional Review Act, Congress subsequently undid this protection. Specifically, Congress passed a joint resolution that provided for the disapproval of the regulation, which President Trump signed into law on April 3, 2017.³² Because this was done under the Congressional Review Act,³³ not only did the resolution kill the FCC’s pending protections, it also prohibited federal agencies from passing any substantively similar regulation under current law.³⁴ As

excluded from the duty to obtain consent. *Id.* Arguably, the most important and contentious of these were the mandate to obtain opt-in consent for sensitive information and the mandate to provide users with an opt-out for non-sensitive information. *See id.*

27. Protecting the Privacy of Customers of Broadband and Other Telecommunications Services, 32 FCC Rcd 1,793, (2017) (order granting stay petition in part).

28. *Id.*

29. Joint Opposition to Petition to Stay Final Rule: Protecting the Privacy of Broadband and Other Telecommunications Services (Feb. 3, 2017), https://consumerfed.org/wp-content/uploads/2017/02/2-3-17-Opposition-to-Stay_Comment.pdf.

30. Protecting the Privacy of Customers of Broadband and Other Telecommunications Services, *supra* note 27, at 3.

31. Protecting the Privacy of Customers of Broadband and Other Telecommunications Services, *supra* note 27, at 3.

32. Act of Apr. 3, 2017, Pub. L. No. 115–22, 131 Stat. 88 (2017) (“Congress disapproves the rule submitted by the Federal Communications Commission relating to ‘Protecting the Privacy of Customers of Broadband and Other Telecommunications Services’ (81 Fed. Reg. 87274 (Dec. 2, 2016)), and such rule shall have no force or effect.”); *see also* Steve Lohr, *Trump Completes Repeal of Online Privacy Protections from Obama Era*, N.Y. TIMES (Apr. 3, 2017), <https://www.nytimes.com/2017/04/03/technology/trump-repeal-online-privacy-protections.html> (reporting on the process through which, “President Trump on Monday signed a congressional resolution to complete the overturning of internet privacy protections created by the Federal Communications Commission”).

33. *See* Congressional Review Act, 5 U.S.C. §§ 801–08 (1996). The Act includes the Congressional Disapproval Procedure, which allows Congress to issue a joint resolution that rescinds a regulation within sixty days of the regulation’s promulgation date. *See* § 802.

34. By May 2017, the 115th Congress had rescinded fourteen regulations pursuant to the Congressional Review Act. *See* Act of Feb. 14, 2017, Pub. L. No. 115–4, 131 Stat. 9 (2017) (rescinding Disclosure of Payments by Resource Extraction Issuers, 81 Fed. Reg. 49359 (July 27, 2016), a rule submitted by the Securities and Exchange Commission); Act of Feb. 16, 2017, Pub. L. No. 115–5, 131 Stat. 10 (2017) (rescinding Stream Protection Rule, 81 Fed. Reg. 93066 (Dec. 20, 2016), submitted by

a result, Congress has effectively barred the FCC from creating other similar consumer-protection rules regulating ISPs. Moreover, because ISPs continue to be classified as common carriers, they are excluded from the FTC Act and remain under the exclusive jurisdiction of the FCC.³⁵ Therefore, the FTC lacks the jurisdiction to impose its own privacy regulations.³⁶

the Office of Surface Mining Reclamation and Enforcement of the Department of the Interior); Act of Feb. 28, 2017, Pub. L. No. 11–8, 131 Stat. 15 (2017) (rescinding Implementation of the NICS Improvement Amendments Act of 2007, 81 Fed. Reg. 91702 (Dec. 19, 2016), submitted by the Social Security Administration); Act of Mar. 27, 2017, Pub. L. No. 115–11, 131 Stat. 75 (2017) (rescinding Federal Acquisition Regulation, 81 Fed. Reg. 58562 (Aug. 25, 2016), submitted by the Department of Defense, the General Services Administration, and the National Aeronautics and Space Administration); Act of Mar. 27, 2017, Pub. L. No. 115–12, 131 Stat. 76 (2017) (rescinding Resource Management Planning, 81 Fed. Reg. 89580 (Dec. 12, 2016), submitted by the Bureau of Land Management of the Department of the Interior); Act of Mar. 27, 2017, Pub. L. No. 115–13, 131 Stat. 77 (2017) (rescinding Elementary and Secondary Education Act of 1965, 81 Fed. Reg. 86076 (Nov. 29, 2016), submitted by the Department of Education); Act of Mar. 27, 2017, Pub. L. No. 115–14, 131 Stat. 78 (2017) (rescinding Teacher Preparation Issues, 81 Fed. Reg. 75494 (Oct. 31, 2016), submitted by the Department of Education); Act of Mar. 31, 2017, Pub. L. No. 115–17, 131 Stat. 81 (2017) (rescinding Federal-State Unemployment Compensation Program; Middle-Class Tax Relief and Job Creation Act of 2012 Provision on Establishing Appropriate Occupations for Drug Testing of Unemployment Compensation Applicants, 81 Fed. Reg. 50298 (Aug. 1, 2016), submitted by the Department of Labor); Act of Apr. 3, 2017, Pub. L. No. 115–20, 131 Stat. 86 (2017) (rescinding Non-Subsistence Take of Wildlife, and Public Participation and Closure Procedures, on National Wildlife Refuges in Alaska, 81 Fed. Reg. 52247 (Aug. 5, 2016), submitted by the Department of the Interior); Act of Apr. 3, 2017, Pub. L. No. 115–21, 131 Stat. 87 (2017) (rescinding Clarification of Employer's Continuing Obligation to Make and Maintain an Accurate Record of Each Recordable Injury and Illness, 81 Fed. Reg. 91792 (Dec. 19, 2016), submitted by the Department of Labor); Act of Apr. 13, 2017, Pub. L. No. 115–24, 131 Stat. 90 (2017) (rescinding Savings Arrangements Established by Qualified State Political Subdivisions for Non-Governmental Employees, 81 Fed. Reg. 92639 (Dec. 20, 2016), submitted by the Department of Labor); Act of May 12, 2017, Pub. L. No. 115–33, 131 Stat. 845 (2017) (rescinding Metropolitan Planning Organization Coordination and Planning Area Reform, 81 Fed. Reg. 93448 (Dec. 20, 2016), submitted by the Federal Highway Administration and the Federal Transit Administration); Act of May 17, 2017, Pub. L. No. 115–35, 131 Stat. 848 (2017) (rescinding Savings Arrangements Established by States for Non-Governmental Employees, 81 Fed. Reg. 59464 (Aug. 30, 2016), submitted by the Department of Labor). Prior to January of 2017, this power had been used only once. *See* Act of Mar. 20, 2001, Pub. L. No. 107–5, 115 Stat. 7 (2001) (disapproving a rule from the Department of Labor relating to ergonomics).

35. *See* Federal Trade Commission Act § 5(a), 15 U.S.C. § 45(a)(2) (2012) (excepting from FTC regulation “common carriers subject to the Acts to regulate commerce”); *see also* Calli Schroeder, *The AT&T v. FTC Common Carrier Ruling Creates a Regulatory “Blind Spot,”* INT’L ASS’N PRIVACY PROFESSIONALS (Sept. 2, 2016), <https://iapp.org/news/a/the-att-v-ftc-common-carrier-ruling-and-how-it-changes-common-carrier-regulation/> (reporting on the implications for privacy of the U.S. Circuit Court ruling that common carriers are exempt from all FTC Act Section 5 actions).

36. However, the FCC can still regulate ISPs through the Communications Act, even similarly to how the FTC regulates internet based companies, as long as these regulations are not deemed substantively similar to the FCC Privacy Order. *See supra* note 33, at § 801(b)(2) (“A rule that does not take effect . . . may not be reissued in *substantially* the same form, and a new rule that is *substantially* the same . . . may not be issued . . .”) (emphasis added).

B. PRIVACY IMPLICATIONS OF ISPS

Most of the central arguments made opposing the FCC Privacy Order were unpersuasive. The four most common of these were that: (i) having two agencies involved in the regulation of digital privacy would create confusion for consumers; (ii) it would stifle innovation; (iii) it would raise costs, which could in turn lead to increased prices for consumers, and; (iv) by restricting consumers' ability to trade their online data for other services, it would reduce consumer choice.³⁷ As it turns out, these arguments largely echo the arguments made a year earlier against net neutrality, which, for the most part, were simply a slight modification of the generic arguments against any regulation.

In the midst of these boilerplate concerns, one argument stands out: That the FCC Privacy Order was unfair for two reasons. First, because it regulated one sector of the behavioral advertising market (ISPs) more than others (websites) and, second, because it created barriers to entry for companies (ISPs) that might have wished to compete with nearly monopolistic incumbents (Alphabet and Facebook).³⁸ While this argument is superficially plausible, it misses the mark. While it is certainly the case that Alphabet, Facebook, and other websites collect vast amounts of personal data, the amount that they can collect is far less than what is available to an ISP. Moreover, even if each ISP has a smaller market share than Alphabet or Facebook in the overall market, for many consumers, ISPs operate as de facto monopolists. Depending on her place of residence, a consumer might be able to choose between only two ISPs or, in some cases, have no choice at all.³⁹

37. See Jeff Flake, *Settling a Bureaucratic Turf War in Online Privacy Rules*, WALL ST. J. (Mar. 1, 2017, 7:06 PM), <https://www.wsj.com/articles/settling-a-bureaucratic-turf-war-in-online-privacy-rules-1488413165> (using equivalent arguments and, on occasion, similar phrasing, to the letter of support from the White House and the letter sent to Congress by the industry coalition); see also Kimberly Kindy, *How Congress Dismantled Federal Internet Privacy Rules*, WASH. POST (May 30, 2017), https://www.washingtonpost.com/politics/how-congress-dismantled-federal-internet-privacy-rules/2017/05/29/7ado6e14-2f5b-11e7-8674-437ddb6e813e_story.html?noredirect=on&utm_term=.d4611d684368 (reviewing the arguments in favor of the measure).

38. See Letter from Jacquelyne Fleming, Asst. Vice President-Fed. Regulatory, AT&T, to Marlene H. Dortch, Sec'y, Fed. Comm'n Comm'n (June 28, 2016), https://ecfsapi.fcc.gov/file/1042634095372/WISPA%20Rebuttal%20Ex%20Parte_Final.pdf ("The proposed marketing restrictions would irrationally protect market incumbents [e.g., Google and Facebook] against competition from new entrants [ISPs] in the digital advertising market.") [hereinafter *Letter to the FCC*]; see also Larry Downes, *Why Congress's Rejection of Proposed FCC Data Rules Will Not Affect Your Privacy in the Slightest*, FORBES (Mar. 30, 2017, 6:00 AM), <https://www.forbes.com/sites/larrydownes/2017/03/30/why-congresss-rejection-of-proposed-fcc-data-rules-will-not-affect-your-privacy-in-the-slightest/>.

39. According to the most recently available data from the FCC, only 43.5% of the US population was serviced by three or more broadband providers offering 25/3 Mbps speed or greater. 33% of Americans had access to services from 2 providers, and a further 19% had access to services from only 1 provider. *Compare Broadband Availability in Different Areas*, FED. COMM'N COMM'N

Further undermining this argument is the fact that under the *status quo*, companies that run websites—such as Alphabet and Facebook—are subject to FTC privacy regulations. Since the Fair Credit Reporting Act was passed in the 1970s, the FTC has been the agency primarily responsible for protecting consumer privacy.⁴⁰ In addition to enforcing statutory privacy laws, the FTC regularly publishes Privacy Reports and works to advance consumer privacy policy in the marketplace. Common carriers, however, fall outside the “marketplace” for regulatory purposes, and are therefore beyond the reach of FTC regulations and enforcement. These entities are instead regulated by their specific agencies, such as the FCC.⁴¹ Because of their unique social and economic role, common carriers are generally subject to more stringent regulations than other private companies.⁴² As a result, one would expect that classifying ISPs as common carriers, and thereby moving them from FTC to FCC jurisdiction, would result in them being subject to more regulation—as are telecommunication companies—compared to their private counterparts regulated by the FTC.

This is not what has happened. Ironically, as a result of Congress’s disapproval,⁴³ ISPs, as common carriers, are now subject to *less* stringent privacy regulations than their non-common carrier counterparts in the digital world. Perhaps more ironic is the fact that ISPs are positioned to collect far more information than their non-common carrier digital counterparts. Unlike, for example, Alphabet, ISPs have many sources of individuals’ personal information. Alphabet can only gather information about an individual consumer when she is using its websites (such as its search engine www.google.com) and other products and services (for example, its web-based email service Gmail or its Chrome web browser). On the other hand, that consumer’s ISP can, in principle, see *everything* that she does on both her computer and her smartphone. It can see every URL she visits, along with every video and song she streams, and every

(last updated Dec. 2016), https://broadbandmap.fcc.gov/#/area-comparison?selectedTech=acfosw&selectedSpeed=25_3.

40. See, e.g., Woodrow Hartzog & Daniel J. Solove, *The Scope and Potential of FTC Data Protection*, 83 GEO. WASH. L. REV. 2230, 2251 (2015); Steven Hetcher, *The FTC as Internet Privacy Norm Entrepreneur*, 53 VAND. L. REV. 2041, 2042 (2000); see also *Fed. Trade Comm’n v. Wyndham Worldwide Corp.*, 10 F. Supp. 3d 602, 611–12 (D.N.J. 2014).

41. See Ohm, *supra* note 4, at 1421 (explaining that, while ISPs lack a regulatory framework, “Congress has already regulated ISP surveillance with the Electronic Communications Privacy Act”). But see Ohm, *supra* note 4, at 1478 (“As many courts and scholars have complained, the ECPA is confusing The rules are particularly confusing for ISP monitoring, because so many exceptions in the law apply to providers, and because courts have had little occasion to consider ISP monitoring.”).

42. See generally Communications Act of 1934, 47 U.S.C. § 151 (1934); see also Rob Frieden, *The FCC’s Name Game: How Shifting Regulatory Classifications Affect Competition*, 19 BERKELEY TECH. L.J. 1275, 1281 (2004).

43. See discussion *supra* Subpart I.A.

file she downloads.⁴⁴ By triangulating cellphone signals, it can also collect data on her location, even if she has disabled her phone's location service.⁴⁵

In other words, at least from a technological standpoint, it is currently feasible for ISPs to collect, and then sell, a staggering amount of data about nearly everyone in the United States.⁴⁶ The limits of these practices depend only on the current FCC's interpretation of Title II of the Communications Act and each ISP's individual privacy policies. While perfectly rational agents might be able to bargain for increased privacy protections if they so desired, individuals affected by NBLLN cannot comprehend how much privacy they are giving up, and are therefore unlikely to bargain for increased protection.

Of course, under the status quo, consumers are in theory still able to opt-out and prevent the ISP from collecting their data in the first place. But, this is easier said than done. Most privacy enhancing technologies, such as a virtual private network ("VPN") or the use of a browser's "private" or "incognito" mode, are reasonably effective at limiting tracking by websites, but are ineffective against tracking by ISPs. The reason for that is that the benefits of these technologies only appear after the consumer has connected to the ISP server. Some technologies—for example, "Tor" (also called "The Onion Router")—are effective at masking a user's identity from the ISP, but they come at the cost of usability and speed.⁴⁷ This form of "opt-out" operates by circumventing

44. See Ohm, *supra* note 4, at 1423 ("An ISP controls a valuable and privileged bottleneck. It owns the point on the network between a user's computer and the rest of the Internet [T]he ISP's connection to the end user, is a unique and critical point: the only point through which all of a user's communications must pass [T]he greatest point of control and surveillance.").

The exception to this, which is the one thing her ISP cannot see, is encrypted communications. This is the case for any site that has an SSL certificate (https sites). When the consumer is in such sites, her ISP knows that she is at the website, but doesn't know what she does in it.

45. For an illustration of how this can produce relevant legal consequences, see *United States v. Carpenter*, No. 12–20218, 2013 WL 6385838 (E.D. Mich. Dec. 6, 2013) (explaining how the government procured 152 days of historical cell phone location data from Timothy Carpenter as key evidence for a criminal investigation). See generally Ohm, *supra* note 4, at 1450 ("Telephone companies and their employees are sued and criminally charged more often than ISPs, usually for installing devices such as pen registers, which record telephone numbers dialed from a phone, and even occasionally for recording voice conversations . . .").

46. Note that, despite this technological feasibility, to date, ISPs do not have a record of violating their users' privacy. See Ohm, *supra* note 4, at 1450 ("No reported cases to date have discussed the liability of an ISP for unlawfully running packet sniffers, except for lawsuits against providers for supporting government monitoring."). However, "[t]here are convincing reasons to suspect that providers have respected privacy only because they have been constrained from doing more[,] . . . technological barriers to extensive monitoring have fallen significantly."). Ohm, *supra* note 4, at 1450.

47. *Tor: Overview*, TOR PROJECT, <https://www.torproject.org/about/overview.html.en> ("Using Tor protects you against a common form of Internet surveillance known as 'traffic analysis.' Traffic analysis can be used to infer who is talking to whom over a public network. Knowing the source and

the ISP's efforts to track the user's online activities. It is only by limiting tracking and, therefore, the amount of information that the ISP is able to collect, that a consumer can meaningfully regulate the ISP's use of that information.

C. THE FCC PRIVACY ORDER AS A DEFAULT RULE

At the heart of this policy dispute is the classic issue of selecting a default rule. Default-based policies rely on the *status quo* bias, which is the idea that when offered a choice containing a default option, people tend to remain in that default.⁴⁸ The effect of the *status quo* is to create a default in which ISPs have virtually unrestricted access to their customers' personal information. The FCC's rules would have changed the default to one in which ISPs would have had to obtain customers' explicit consent before obtaining, and by extension using, their personal information. This would have shifted the default from an opt-out rule to an opt-in rule.

This change would have effectively reversed the interpretation of silence on the part of the customer. Under the FCC's disallowed rules, in order to use a customer's personal information, ISPs would have first been required to obtain that customer's explicit consent.⁴⁹

Default rules designed to nudge may be either "policy defaults" or "penalty defaults." Policy defaults aim to increase the number of people choosing the default option. Penalty defaults, introduced by Ayres and Gertner,⁵⁰ aim to encourage a private party to provide information to other parties, thereby reducing rent-seeking (obtaining gains by generating uncompensated losses to others) under information asymmetries.⁵¹

Superficially, the choice of a default rule may seem relatively unimportant. Unlike a mandatory rule, a default rule is simply a gap-filler, which provides the rules in the event that parties do not supply their own.⁵² There is, however, a long line of literature on the importance of default rules. Default rules are known to be "sticky," meaning that

destination of your Internet traffic allows others to track your behavior and interests.") (last visited July 29, 2018).

48. See Charles J. Goetz & Robert E. Scott, *The Limits of Expanded Choice: An Analysis of the Interactions Between Express and Implied Contract Terms*, 73 CALIF. L. REV. 261, 321 (1985).

49. See *FCC Privacy Order*, *supra* note 17, at 2506.

50. Ian Ayres & Robert Gertner, *Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules*, 99 YALE L.J. 87, 91 (1989).

51. See *id.* at 94; see also Ian Ayres & Robert Gertner, *Strategic Contractual Inefficiency and the Optimal Choice of Legal Rules*, 101 YALE L.J. 729, 736 (1992).

52. See, e.g., Ayres & Gertner, *supra* note 50.

parties contract around them far less often than scholars would expect.⁵³ Moreover, law sometimes makes them artificially “stickier” when it defines the way in which parties can contract around the default legal treatment.⁵⁴

The importance of “sticky” defaults may be one reason why an ISP could, in theory, change its terms of service to restrict its use of this information. However, to our knowledge, none has.⁵⁵ We argue that this phenomenon is compounded by the well-documented fact that—due to NBLLN—individuals tend to underestimate their privacy loss in exactly these sorts of contexts.⁵⁶

The FCC Privacy Order can be understood as a set of penalty default rules.⁵⁷ In the absence of a penalty default or other disclosure rule, information asymmetries can lead to situations where the more informed counterparty withholds information from her less-informed counterpart. Such a reduction is socially inefficient, in the sense that it reduces the total value (or “surplus”) created by the exchange. As such, by

53. See, e.g., Colin F. Camerer, *Prospect Theory in the Wild: Evidence from the Field*, in CHOICES, VALUES AND FRAMES 294 (Daniel Kahneman & Amos Tversky eds., 2000); David Cohen & Jack L. Knetsch, *Judicial Choice and Disparities Between Measures of Economic Values*, 30 OSGOOD HALL L.J. 737 (1992); Julie S. Downs et al., *Strategies for Promoting Healthier Food Choices*, 99 AM. ECON. REV. 159 (2009) (illustrating the power of defaults by showing that the effect is also present in insurance, food choices, and marketing, where the number of consumers who agree to receive marketing e-mails increases up to 50% depending on the default); Eric J. Johnson et al., *Defaults, Framing and Privacy: Why Opting In-Opting Out*, 13 MARKETING LETTERS 5 (2002); Eric J. Johnson & Daniel G. Goldstein, *Defaults and Donation Decisions*, 78 TRANSPLANTATION 1713, 1714–15 (2004) (showing that sticky defaults exist even when the costs of switching away from the default choice are close to zero); Eric Johnson & Daniel Goldstein, *Do Defaults Save Lives*, 302 SCIENCE 1338–1339 (2013) (illustrating the power of defaults by showing that the number of organ donors increases up to 400% in countries where being a donor is the default choice); Daniel Kahneman et al., *Anomalies: The Endowment Effect, Loss Aversion, and Status Quo Bias*, 5 J. ECON. PERSP. 193, 197, 199 (1991); Brigitte Madrian & Dennis Shea, *The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior*, 116(4) Q. J. ECON. 1149, 1152 (2001) (illustrating the power of defaults by showing that the variety of domains in which a status-quo bias exists: adherence to savings plans increases up to 50% when employees are enrolled automatically); William Samuelson & Richard Zeckhauser, *Status Quo Bias in Decision Making*, 1 J. RISK & UNCERTAINTY 7, 44 (1988); Cass R. Sunstein, *Impersonal Default Rules vs. Active Choices vs. Personalized Default Rules: A Triptych* (May 19, 2013) (unpublished manuscript) (available at https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2171343).

54. See Ian Ayres, *Regulating Opt-Out: An Economic Theory of Altering Rules*, 121 YALE L.J. 2032, 2084 (2012) (providing a theory of altering rules and discussing “impeding altering rules,” which deter some parties from choosing legally disfavored provisions).

55. Under the status quo, customers are still, in theory, able to prevent an ISP from using that information by preventing the ISP’s efforts to collect the data in the first place. However, as discussed in Subpart I.B. above, this is easier said than done.

56. See *infra* Part III.

57. See Ignacio N. Cofone, *The Way the Cookie Crumbles: Online Tracking Meets Behavioural Economics*, 25 INT’L J.L. & INFO. TECH. 38, 48 (2016); Jay P. Kesan & Rajiv C. Shah, *Setting Software Defaults: Perspectives from Law, Computer Science and Behavioral Economics*, 82 NOTRE DAME L. REV. 583, 620–21 (2006) (analyzing the similar case of online tracking through http cookies in the context of default rules).

withholding information, the more informed party is essentially destroying value. Despite this value destruction, it is rational for the more informed party to do so. While the total surplus—the benefit from the transaction—is reduced, the amount of the surplus taken from the less informed party is larger than the more informed party's proportional share of the destroyed surplus. When this happens, a rule encouraging the party to reveal the information can prevent this self-interested but destructive behavior, increasing the total social value. Penalty defaults are intended to do just this and, therefore, can sometimes be used to counteract incentives to strategically withhold information, just as presumptions can do so in procedural law.⁵⁸

D. ECONOMIC ARGUMENTS FOR INTERVENTION

Standard economic theory argues that competition leads to efficient outcomes,⁵⁹ and that the best thing a government or regulator can do is to stay out of the way, and at most, help make lump-sum cash transfers between individuals.⁶⁰ While this argument has lost some of its luster in recent decades, economists are generally sympathetic towards the virtues of markets. In most markets, competition does a good job of ensuring that the needs of consumers are met. There is also a large literature pointing out the fact that regulations can have unintended consequences, and can ultimately harm consumers.⁶¹

58. See Ayres & Gertner, *supra* note 50, at 107; see also Cofone, *supra* note 57, at 48–49.

59. For example, the First Fundamental Theorem of Welfare Economics (sometimes called simply the “First Welfare Theorem”) asserts that, under certain conditions, “competitive equilibrium allocations are Pareto optimal,” meaning that there is no other allocation that can make all parties better off. TRUMAN F. BEWLEY, *GENERAL EQUILIBRIUM, OVERLAPPING GENERATIONS MODELS, AND OPTIMAL GROWTH THEORY* 17, 160–61 (2008). For a proof of the theorem, see *id.* at 162–63; see also ANDREU MAS-COLELL ET AL., *MICROECONOMIC THEORY* 325–327 (1995).

60. For example, the Second Fundamental Theorem of Welfare Economics (sometimes called simply the “Second Welfare Theorem”) states that, under certain conditions, “[A]ny Pareto optimal allocation can be achieved as the allocation of a competitive equilibrium after an appropriate lump-sum redistribution of wealth among consumers.” BEWLEY, *supra* note 59, at 160. For a proof of the theorem, see BEWLEY, *supra* note 59, at 172–76; see also ANDREU MAS-COLELL ET AL., *supra* note 59, at 327–28.

61. See, e.g., Cass R. Sunstein, *Political Equality and Unintended Consequences*, 94 COLUM. L. REV. 1390 (1994); see also Feng Gao et al., *Unintended Consequences of Granting Small Firms Exemptions from Securities Regulation: Evidence from the Sarbanes-Oxley Act*, 47 J. ACCT. RES. 459 (2009) (finding that exemptions for small firms in the Sarbanes-Oxley Act of 2002 have had the unintended effect of encouraging small firms to stay small); Robert E. Litan & Hal J. Singer, *Unintended Consequences of Net Neutrality Regulation*, 5 J. TELECOMM. & HIGH TECH. L. 533 (2006) (concluding that certain net neutrality requirements had the unintended consequence of reducing innovation); Ekaterina Jardim et al., *Minimum Wage Increases, Wages, and Low-Wage Employment: Evidence from Seattle* (Nat'l Bureau of Econ. Research, Working Paper No. 23532, 2017), <http://www.nber.org/papers/w23532> (finding evidence that a minimum wage increase in Seattle resulted in a reduction in both employment and total payroll in low wage jobs).

In this standard framework, arguments for market intervention typically rely on the existence of a market failure. One traditional market failure is in the realm of public goods. These are goods or services that are both non-rivalrous in consumption (consumption by one person does not affect its value to another) and non-excludable (it is very difficult or impossible to exclude individuals from the benefits of the good or service once it is being provided), such as national defense.⁶² Another traditional cause of market failures is the existence of monopolies, which provides a basis for antitrust regulation.⁶³

More recently, scholars have begun to point to consumer irrationality as another potential justification for market intervention.⁶⁴ While some have expressed skepticism about the increasing prominence of behavioral economics,⁶⁵ it is now widely accepted in legal scholarship that, under the right circumstances, the existence of a sufficiently prevalent cognitive bias can be grounds for market intervention. In Part III, we argue that this is the case in the context of internet data privacy.

E. THE PRIVACY PARADOX AND NBLLN

One of the complications in the discussion of consumer data privacy is known as the privacy paradox. This paradox refers to the fact that, while individuals say that they are concerned about their privacy, they are willing to sell or trade this same privacy for almost nothing.⁶⁶

One response to this paradox is skepticism. When there is a conflict between what an individual reports and what she does, it may be prudent to put more weight on what the individual actually does. Statements may reflect aspirations, while actions may be more likely to reflect tough tradeoffs made in the face of real life costs and benefits. This is a major reason why economists and psychologists tend to view actions as much more informative than words.

This logic of rational decision-making for privacy choices breaks down in the face of a systematic cognitive bias. If individuals deviate from

62. ROBERT COOTER & THOMAS ULEN, *LAW AND ECONOMICS* 45–46 (5th ed. 2007).

63. *See id.* at 43.

64. *See, e.g.*, Cass R. Sunstein, *The Storrs Lectures: Behavioral Economics and Paternalism*, 122 *YALE L.J.* 1826 (2013); Christine Jolls & Cass R. Sunstein, *Debiasing Through Law*, 35 *J. LEGAL STUD.* 199 (2006); Colin Camerer et al., *Regulation for Conservatives: Behavioral Economics and the Case for “Asymmetric Paternalism,”* 151 *U. PA. L. REV.* 1211 (2003); Cass R. Sunstein & Richard H. Thaler, *Libertarian Paternalism Is Not an Oxymoron*, 70 *U. CHI. L. REV.* 1159 (2003); Cass R. Sunstein, *Switching the Default Rule*, 77 *N.Y.U. L. REV.* 106 (2002); Christine Jolls et al., *A Behavioral Approach to Law and Economics*, 50 *STAN. L. REV.* 1471 (1998); Sunstein, *supra* note 53.

65. *See infra* Part II.B.2.

66. *See, e.g.*, H. Brian Holland, *Privacy Paradox 2.0*, 19 *WIDENER L.J.* 893, 893 (2010); Norberg et al., *supra* note 16, at 100; Eric Jorstad, *The Privacy Paradox*, 27 *WM. MITCHELL L. REV.* 1503, 1503 (2001) (“Americans are ambivalent about privacy.”).

standard notions of rationality—for example, because of NBLN—they may *think* that their actions are perfectly consistent with their statements. The cognitive bias flips the normal situation on its head: it is the statement, rather than the action, that is a true reflection of the individual's preferences.

Rather than dismissing the possibility that consumers act irrationally when making privacy choices, we consider the possibility that people making online choices about their personal information might face this widely recognized cognitive bias. We then explore the implications of this extended decision-making model, and discuss how the law can help protect individuals and increase social welfare.

II. PRIVACY LOSS IN THE FACE OF NBLN

A. OUR MODEL OF PRIVACY LOSS

1. *Types and Clues*

In prior work, we presented a model to formalize the concept of privacy loss based on Bayesian updating.⁶⁷ Here, we consider a simplified version of that model, which conveys its main thrust. Consider an individual named Abby. Abby has a fundamental characteristic. This might be her height, her willingness to pay for a good, her wealth, her desirability as an employee, or her intrinsic worth as human being. We will refer to this as her “type.” Initially only Abby knows her type.

Now consider a company called Poodle. Suppose that initially, Poodle has no specific information about Abby, but would like to learn more about her. While Poodle cannot observe Abby's type directly, it does have a pretty good idea about what the overall distribution of types in the population looks like. In addition to the mean and the standard deviation, it also knows the general shape of the distribution—for example, whether individuals are pretty evenly spread out across different types, or whether they tend to be bunched together with only a few outliers.

Poodle can also observe signals, or clues, that allow it to guess something about Abby's type. Each of these clues represents a piece of information about Abby. While none of these clues fully reveal Abby's type on their own, by running analytics on these clues, Poodle can form a clearer picture about it. Specifically, when it aggregates these clues, Poodle can form its best guess about Abby's type. Because it knows that this is only an informed guess, it still has some uncertainty about her type—it might guess too high or too low, for example, believing her to be

67. Cofone & Robertson, *supra* note 14.

taller or shorter than she actually is. The more uncertain Poodle is about Abby's type, the more privacy she has.

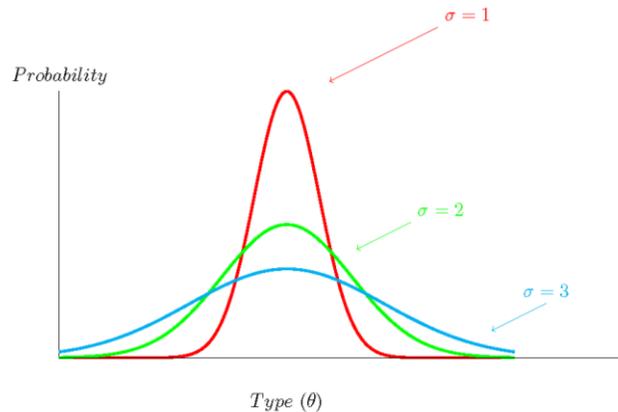


Figure 1: Abby's privacy is higher when the distribution has fatter tails

Figure 1 illustrates this intuition. When Poodle has few signals about Abby's type ($\sigma=3$, the widest, blue curve), it knows that the range of plausible types is wide. As it gets to know Abby better (moving from the widest curve to the intermediate, green curve, $\sigma=2$), the distribution becomes narrower, meaning that the range of plausible types narrows. As it gets to know Abby better still (moving to the narrowest, red curve, $\sigma=1$), Poodle has a good idea about what it wants to know about Abby, and the range of plausible values becomes narrower.

2. Aggregating Clues

It turns out that, when faced with large amounts of data, the average person is not particularly good at estimating the informativeness of each piece of newly arriving information.⁶⁸ 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. That is, they might mistakenly believe that Poodle's belief about their personal information does not change much when they provide the company with new information that is, in fact, still informative. In short, they will

68. See Daniel Kahneman & Amos Tversky, *Subjective Probability: A Judgment of Representativeness*, 3 *COGNITIVE 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").

(mistakenly) think that Poodle cannot learn much from the new information.⁶⁹

The upshot of this is that individuals will tend to give away their private data too easily, or sell it too cheaply. An example can help to clarify this phenomenon. Suppose that 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 the company that created the application, Poodle, 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 it is combined with the data that she shares, through a different application, with a second company, called Goggles. In other words, she will underestimate her privacy loss, causing her to undervalue her private data.⁷⁰ 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 other words, she may accept terms that, but for her misunderstanding, she would not accept.

While the effect of NBLLN on an individual's 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. These individuals 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.

69. This relates to Abby's perceived number of draws. For example, Abby might give Poodle ten extra pieces of information (ten signals), mistakenly believing that their effect on Poodle's posterior is the same as (or similar to) the effect of giving it three extra pieces of information. In that case, Poodle's posterior would be tighter than what Abby believes it to be. If Abby has decreasing marginal utility over her privacy, then, she will systematically underestimate her privacy harm for each subsequent signal. We address this in more detail below. *See infra* note 87.

70. It is also possible that Abby misperceives Poodle's prior. We address this possibility later in the analysis. *See infra* note 87.

B. NON-BELIEF IN THE LAW OF LARGE NUMBERS

1. *A Well-Known Bias*

Our baseline model of privacy loss was built using standard statistical principles. In doing so, we implicitly assumed that individuals are rational Bayesians: they begin with some set of beliefs, and then update their beliefs based on things that they see in the world. We also assumed that the more information they see, the more updating they do. At the extreme, if Poodle observed an infinite number of pieces of information about Abby, Poodle would know Abby's type with absolute certainty. This fundamental statistical principle is known as the Law of Large Numbers.

It turns out that most humans have very poor intuition when it comes to fundamental statistical principles. One place where this is particularly true is with respect to the Law of Large Numbers. For example, in a famous paper, Kahneman and Tversky demonstrated that individuals systematically underestimate how quickly a sample mean converges to the population mean.⁷¹ In one setting, they asked respondents the following question, which has since become a classic formulation:

A certain town is served by two hospitals. In the larger hospital about 45 babies are born each day and in the smaller hospital about 15 babies are born each day. As you know, about 50% of all babies are boys. The exact percentage of baby boys, however, varies from day to day. Sometimes it may be higher than 50%, sometimes lower.

For a period of 1 year, each hospital recorded the days on which (more/less) than 60% of the babies born were boys. Which hospital do you think recorded more such days?⁷²

Subjects were given three answer options: (i) the larger hospital, (ii) the smaller hospital, and (iii) about the same (i.e., within 5% of each other). Respondents were divided into two groups: those who were asked about days on which more than 60% of babies born were boys, and those who were asked about days on which less than 60% of babies born were boys.

Statistically speaking, there is a clear answer to the two questions: the larger hospital is much more likely to have recorded less than 60% boys, while the smaller hospital is much more likely to have recorded more than 60% boys. This is because, as a sample increases, the likelihood that the sample mean (in the example, the percentage of babies that are boys) will diverge very much from the true population parameter (here, 50%) drops fairly rapidly. Moreover, because of the statistical

71. See Kahneman & Tversky, *supra* note 68, at 437–445.

72. See Kahneman & Tversky, *supra* note 68, at 443.

phenomenon known as the Law of Large Numbers, as the sample size approaches infinity, the probability that the sample mean will diverge from the true parameter *at all* falls to zero.⁷³

While this statistical fact is well established, applying it requires considerable mental effort, and does not appear to come naturally to individuals. In Kahneman and Tversky's original study, the majority of respondents incorrectly answered that the probabilities were about the same, despite the fact that the larger hospital is three times the size of the smaller one. In fact, even more worrisome, for both groups, the correct answer was actually the *least* popular choice.

Benjamin et al. summarize Kahneman and Tversky's findings as evidence that, "experimental subjects seem to think sample proportions reflect a 'universal sampling distribution,' virtually neglecting sample size."⁷⁴ As a result, these individuals systematically over-estimate the level of uncertainty—this is, the standard deviation of the sampling distribution—for large samples. In the context of online data privacy, this bias implies that an individual like Abby will tend to underestimate how much a company like Poodle can learn about her by analyzing her online behavior, leading her to undervalue her private data. Benjamin et al. go on to develop a relatively tractable⁷⁵ mathematical model of NBLLN.⁷⁶ We adopt this framework for our analysis of the effects of NBLLN on privacy loss. We are the first to apply this concept to the privacy context.

2. *Addressing Concerns About Behavioral Economics*

A common critique of behavioral economics is that psychologists have documented a great many cognitive biases, many of which ought to act in conflicting directions. As a result, critics argue, it is often difficult to determine both which biases are most important in any particular context, and how important any particular bias ought to be.⁷⁷ We also recognize the concern that researchers may be tempted to rummage around in a psychology textbook until they find a bias that suits their story.

While we recognize and share these concerns, neither of these critiques are compelling in this context. First, there is strong evidence

73. See ROBERT V. HOGG ET AL., INTRODUCTION TO MATHEMATICAL STATISTICS 204 (6th ed. 2005) (discussing the law of large numbers).

74. Daniel J. Benjamin et al., *A Model of Nonbelief in the Law of Large Numbers*, 14 J. EUR. ECON. ASS'N 515, 516 (2016).

75. This tractability is very important in attempting to model NBLLN. Precisely because it is so fundamental to statistical reasoning, it is rather difficult to construct a model *without* relying on the Law of Large Numbers. This in turn makes it difficult to model an individual who suffers from NBLLN and explains the importance of the work of Benjamin et al.

76. Benjamin et al., *supra* note 74.

77. See, e.g., Alan Schwartz, *Regulating for Rationality*, 67 STAN. L. REV. 1373 (2015).

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 U.S. adults, Stango et al. studied the incidence of seventeen different widely recognized “behavioral factors,” including NBLLN.⁷⁸ They report that 87% of the participants in their study exhibited NBLLN.⁷⁹ This suggests that a huge proportion of the population is vulnerable to this bias, which should reassure any skeptical reader that NBLLN is not a fringe issue. Rather, it appears to be the dominant way that individuals think about how information aggregates. There is therefore a solid basis for believing that NBLLN plays a substantial role in the realm of online consumer data privacy, and that it could also be driving the well-known privacy paradox.

Moreover, our prior model is *explicitly* premised upon the individual’s ability to judge the degree to which the standard deviation of the probability distribution of someone’s belief over target person’s type shrinks around its mean value.⁸⁰ This may be the most important assumption of the model. NBLLN *explicitly* undermines this assumption in a way that is particularly relevant in a context in which people release many small pieces of information—precisely the context of privacy—relevant internet interactions. It is therefore important from a conceptual perspective to explore the consequences of this bias.

C. NBLLN AND PRIVACY LOSS

1. Formalization

We begin by formalizing the application of Benjamin et al.’s model of NBLLN⁸¹ to our model of privacy loss.⁸² This provides the formal mathematical justification for our argument. We then discuss the intuition behind this formalization.

Suppose that instead of Abby, we have two individuals named Tommy and Sam. Sam is Tommy’s irrational alter ego. The two are identical in every way but one. Tommy is a Bayesian—he correctly

78. Victor Stango et al., *The Quest for Parsimony in Behavioral Economics: New Methods and Evidence on Three Fronts* (Nat’l Bureau of Econ. Research, Working Paper No. 23057, 2017), <http://www.nber.org/papers/w23057>.

79. *See id.* Table 4. This makes NBLLN the second most prevalent behavioral factor from among these seventeen factors, after violation of the general axiom of revealed preference (“GARP”). *Id.*

80. *See* Figure 1.

81. Benjamin et al., *supra* note 74.

82. Cofone & Robertson, *supra* note 14. In contrast to that paper, where we defined privacy loss in terms of standard deviations, here we use the variance of the observers sampling distribution. This is done solely for expositional convenience. Variance is simply standard deviation squared, and working directly with variances allows us to omit additional notation. Moreover, because neither can be negative, there is a one-to-one correspondence between the two concepts.

perceives Poodle's posterior distribution. Sam, on the other hand, suffers from NBLLN.

In particular, suppose that individuals can be represented by a type θ on the interval $(0,1)$. Poodle is able to observe binary signals about Tommy and Sam, which take a value of either y or n . These signals could represent the person's decision about whether or not to stream a particular song, download a particular file, or click on a particular link. The signals are independent and identically distributed (iid) Bernulli draws, where the probability of observing y is θ , leading to a " θ -binomial" distribution of the mean. Tommy understands this, just as he understands that the variance of this distribution decreases at a rate of $1/N$. As N gets very large, Tommy understands that this variance goes to zero. Sam, on the other hand, believes (falsely) that this variance *never* goes to zero. In particular, following Benjamin et al.,⁸³ Sam believes that as N gets large, the distribution will converge to a " β -binomial distribution" for some $\beta \in [0,1]$ that is itself drawn from a distribution with mean θ . Denote this distribution, called Sam's "subjective rate," by $f_{\beta|\theta}^S$.

Benjamin et al. prove that, under relatively mild assumptions,⁸⁴ Sam and Tommy will both perceive the mean of Poodle's distribution correctly. While Sam correctly understands that the variance of Poodle's posterior is decreasing in N , for any $N > 1$ he believes that Poodle's distribution is wider than it really is.

Formally, Tommy understands that, for a sample size N , Poodle's posterior variance is given by

$$\sigma_N^{2T} = \frac{\sigma^2}{N} = \frac{\theta(1-\theta)}{N}.$$

However, Sam incorrectly believes that the posterior variance is given by

$$\sigma_N^{2S} = \frac{\sigma^2}{N} + \frac{N-1}{N} \text{Var}[f_{\beta|\theta}^S] = \frac{\theta(1-\theta)}{N} + \frac{N-1}{N} \text{Var}[f_{\beta|\theta}^S].$$

For the remainder this is analysis, we will assume that $\text{Var}[f_{\beta|\theta}^S]$ is independent of N .

Now suppose that Sam and Tommy are both considering "selling" a clump of N signals to Poodle for a fixed price per signal. How many

83. See generally Benjamin et al., *supra* note 74.

84. See Benjamin et al., *supra* note 74, at 521-22 (noting the model's assumptions).

signals will each be prepared to sell at that price? Standard economic theory indicates that they will be prepared to sell signals up to the point where their marginal disutility from giving up the privacy associated with the signals is equal to the payment offered on the last signal.

We begin by computing Tommy's marginal utility loss from giving up a signal. Since utility is defined over the variance of Poodle's distribution, and this variance is a function of N , we apply the chain rule⁸⁵ to find that

$$\frac{\partial U(\sigma_N^{2T})}{\partial N} = \frac{\partial U(\sigma_N^{2T})}{\partial \sigma_N^{2T}} \frac{\partial \sigma_N^{2T}}{\partial N} = \frac{\partial U(\sigma_N^{2T})}{\partial \sigma_N^{2T}} \frac{-\theta(1-\theta)}{N^2}.$$

Sam's marginal utility loss, in contrast, is given by

$$\frac{\partial U(\sigma_N^{2S})}{\partial N} = \frac{\partial U(\sigma_N^{2S})}{\partial \sigma_N^{2S}} \frac{\partial \sigma_N^{2S}}{\partial N} = \frac{\partial U(\sigma_N^{2S})}{\partial \sigma_N^{2S}} \frac{-\theta(1-\theta) + \text{Var}[f_{\beta|\theta}^S]}{N^2},$$

where $\text{Var}[f_{\beta|\theta}^S] < \theta(1-\theta)$.⁸⁶

Now, suppose that Sam is beginning from the "correct" starting point. Even though he suffers from NBLLN, his beliefs have "caught up," in the sense that he does understand how much Poodle knows about him. In that case,

$$\frac{\partial U(\sigma_N^{2T})}{\partial \sigma_N^{2T}} = \frac{\partial U(\sigma_N^{2S})}{\partial \sigma_N^{2S}} = \frac{\partial U(\sigma_N^2)}{\partial \sigma_N^2}.$$

Since $\text{Var}[f_{\beta|\theta}^S] > 0$, and $\frac{\partial U(\sigma_N^2)}{\partial \sigma_N^2} > 0$, it follows that the absolute value of Sam's marginal utility loss is smaller (closer to zero) than Tommy's. Sam is therefore willing to sell the signals more cheaply than Tommy is or, alternatively, is willing to sell more signals than Tommy is for a given price B .⁸⁷

85. The chain rule is a method for finding the derivative of composite functions and is a fundamental tool of calculus. See GILBERT STRANG, *CALCULUS* 154–56 (1991) (discussing the chain rule).

86. See generally Benjamin et al., *Appendix to: "A Model of Non-Belief in the Law of Large Numbers"* (Mar. 23, 2014), <https://scholar.harvard.edu/files/rabin/files/barney2014.pdf> (presenting the proof).

87. Things get more complicated if Sam hasn't yet realized the true variance of the distribution. In this case, $\sigma_N^{2S} > \sigma_N^{2T}$. If we assume that utility is concave, this implies that

2. Intuition

While the mathematics behind our formalization draw on fairly advanced statistical concepts, the intuition is very simple. Like Kahneman and Tversky's experimental subjects, Sam doesn't make the connection between the size of a pool of data and the precision of the estimates that can be made based on that pool. While he recognizes that the more data you have to look at, the more accurate your estimates based on that data will be, he fails to recognize how quickly the level of precision of these estimates increases. When Sam downloads Poodle's app and accepts its terms of service, he might realize that he is granting Poodle access to his geolocation data. What he may not realize is just how much Poodle can learn about him from that information alone. For example, the app could record the fact that, every weekday morning, he travels at walking pace from 2nd Avenue and 88th Street over to 86th and Lexington before moving at high speed to 59th and Lexington. He then moves at walking speed to 59th and Fifth. Every weekday evening, he reverses the trip. Based on this alone, Poodle can infer where Sam lives and works, and that he commutes by subway. Similarly, his geo-location data between the hours of 11:30 am and 1:30 pm will reveal his lunch routine, just as his location between 6:30 pm and 9:00 pm reveal where he likes to eat dinner. Along the way, the app will collect data that it can use to learn about his shopping habits, his hobbies, and who he socializes with.

In our example, Sam might be horrified to learn this. While he was willing to accept the terms of service, he did so only because he did not

$$\frac{\partial U(\sigma_N^{2T})}{\partial \sigma_N^{2T}} > \frac{\partial U(\sigma_N^{2S})}{\partial \sigma_N^{2S}}.$$

This exacerbates the problem: not only does Sam misperceive the effect of an additional signal on the variance of Poodle's posterior, he *also* misperceives where he is in his own utility function. At the same time, however, this is a countervailing effect—if Sam really misperceives the starting point of Poodle's posterior, he will also *overestimate* the impact of a signal on Poodle's posterior. This follows from the fact that

$$\frac{\partial^2 (\sigma_N^{2S})}{\partial N^2} = \frac{2[\theta(1-\theta) - \text{Var}[f_{\beta|\theta}^S]]}{N^3} > 0.$$

In other words, while σ_N^{2S} is decreasing in N, it does so at a decreasing rate.

As a result, the marginal effect of an incremental signal is larger when the variance is larger. If Sam *thinks* the variance of Poodle's distribution is larger than it really is, he might actually *overestimate* the effect of an additional signal. In the abstract, it is impossible to know which effect will dominate. This generates two related problems. First, Sam, and others like him, is underpricing his private information. Second, while Sam does not realize immediately how much privacy he has left, there is a good chance that, eventually, he will. When he does, he will suffer a severe loss. At the same time, however, the fact that Sam begins by overestimating the variance of Poodle's posterior means that he will overestimate the degree to which an individual signal will affect his privacy. This third effect might offset the first two, leading to indeterminate outcomes.

realize what he was accepting. His lack of understanding, moreover, is driven not by the fact that he did not read the policy, or even that he did not understand the words being used. Rather, the cause of his misunderstanding is that he systematically underestimates how much information can be gleaned from a given set of clues.

III. ADDRESSING BIASED BELIEFS AND INFORMATION OVERLOAD

A. IMPLICATIONS OF THE MODEL

This analysis does more than provide an explanation for the privacy paradox. It also indicates a failure in the market for consumer privacy and provides an economic rationale for regulating consumer privacy. The more consumers need to aggregate information in order to gauge the value of the personal data that they are releasing, the more relevant this rationale will be. Because it is hard to imagine a context in which this is more relevant than that of ISP tracking, it is ironic that ISP tracking is one of the least regulated consumer privacy interactions.

Under normal conditions, standard economic theory suggests that trades are welfare enhancing. If privacy were like fruit, this conclusion would also apply.⁸⁸ Under normal conditions, if Abby has an apple and Poodle has a banana, Abby will only agree to trade with Poodle if she values the banana more than she values the apple. Otherwise, short of coercion or deceit, she will decline to trade. This simple but powerful idea can also be applied to more complex situations. For example, instead of apples and bananas, Abby might have private information about herself, and Poodle might be offering her the right to use its app. Unfortunately, cognitive biases such as NBLLN complicate matters further.

If Abby suffers from NBLLN (as Sam did in Part II.C.), this effectively means that she does not realize how much she values her own information. Going back to fruit, it is almost as if she *thought* she was only trading one apple for one banana, when in fact she was trading a whole bushel of apples. The risk is clear. Abby may be agreeing to a trade that, had she fully understood the situation, she likely would not have agreed to.

Moreover, if there are lots of people like Abby in the world—people who have a hard time distinguishing between a single apple and a bushel of them—it is easy to see how we might end up with a whole lot of apples

88. Of course, there are reasons to believe that privacy is not like fruit. *See* Strandburg, *supra* note 15, at 95 (explaining why the release of personal data in exchange for goods and services is not a typical market and arguing that, “In a functioning market, payment of a given price signals consumer demand for particular good and services, transmitting consumer preferences to producers. Data collection would serve as ‘payment’ in that critical sense only if its transfer from users to collectors adequately signaled user preferences for online goods and services”).

being sold by the Abbys of the world to the Poodles of the world. The Abbys of the world would end up selling off far more apples than they meant to sell, and far more apples than *would have been sold* in an efficient market with fully rational participants. This is a key problem for consumer privacy in a behavioral world.

There are two ways in which the law can address this problem. The first is to tackle it through private law, and particularly the law of contracts. The second is to tackle it through direct regulation—like the FCC attempted to do. We evaluate both of these possibilities.

B. PRIVATE LAW APPROACHES

While we have already shown that *free* contracting will lead to inefficient outcomes, one way to approach this problem is through contract law principles. Contract law has devised ways to address behavioral biases in standard form contracts, chiefly through the doctrine of unconscionability.⁸⁹ Unconscionability has traditionally been divided into two parts: procedural and substantive.⁹⁰ Procedural unconscionability deals with defects in bargaining or contract formation process in a way that is more flexible than other doctrines such as duress, fraud, or incapacity.⁹¹ Substantive unconscionability, on the other hand, allows a judge to void an otherwise valid contract based solely on the fact that the terms of the contract are unfair or oppressive.⁹² While these two concepts are distinct, they are often discussed in tandem.

Russell Korobkin has criticized the modern doctrine of unconscionability as insufficient to the task of dealing with the effect of certain behavioral biases—grouped under the umbrella of bounded rationality—in the context of standard form contracts.⁹³ Korobkin's critiques can also be applied in the context of NBLLN consumers and data privacy.⁹⁴ Moreover, Alan Schwartz has argued persuasively that cognitive errors should be irrelevant to an unconscionability finding, and

89. See generally Russell Korobkin, *Bounded Rationality, Standard Form Contracts, and Unconscionability*, 70 U. CHI. L. REV. 1203 (2003); Jason Scott Johnston, *The Return of Bargain: An Economic Theory of How Standard-Form Contracts Enable Cooperative Negotiation Between Businesses and Consumers*, 104 MICH. L. REV. 857 (2006); Shmuel I. Becher, *Behavioral Science and Consumer Standard Form Contracts*, 68 LA. L. REV. 117 (2007).

90. See Arthur Allen Leff, *Unconscionability and the Code—The Emperor's New Clause*, 115 U. PA. L. REV. 485, 485–489 (1967) (presenting the classic treatment of the unconscionability doctrine).

91. ROBERT E. SCOTT & JODY S. KRAUS, *CONTRACT LAW AND THEORY* 507 (5th ed. 2013).

92. *Id.*

93. Korobkin, *supra* note 89, at 1255–78.

94. While “bounded rationality” is a general term used to describe individuals who do not perfectly process all available information, NBLLN is a much more precise cognitive bias. We can therefore interpret NBLLN as one particular type of bounded rationality.

suggests that courts are particularly poorly positioned to evaluate issues related to cognitive biases.⁹⁵

The issue of institutional competence is particularly concerning in the context of consumer data privacy involving NBLLN. Not only are judges poorly situated to evaluate whether disclosures are sufficiently clear and comprehensible to an NBLLN consumer, an analysis of the substantive terms of the contract would be infeasible. In order for a judge to make a finding of substantive unconscionability in this context, she would have to determine, for each individual consumer, that the amount of personal data being traded was “too large” relative to the benefit that the consumer received. This is a herculean task. After all, as we have discussed, the value of each piece of information about Abby depends crucially on the amount of *other* information that is already known about her.

Indeed, in order to use substantive unconscionability to solve the problems created by NBLLN, a judge would have to make an individualized determination for each affected consumer, in light of all the *other* information already available. Given the number of individuals affected by each standard form data privacy agreement, such an individualized analysis is impractical. Perhaps more importantly, given the prevalence of NBLLN,⁹⁶ there is little reason to believe that a judge would be any better at evaluating the true extent of Abby’s privacy loss than she is.

Finally, even if these issues could be addressed, there remains another, larger problem that is unique to the context of data privacy: once a consumer’s data has been used to learn more about that person, it is virtually impossible to force the user to “unlearn” it. Unconscionability, which is concerned with nullifying a contract *ex post*, is therefore poorly suited to address contracts when the concern is about privacy harms.

An alternative, albeit related, private law approach to this issue is through the doctrine of unilateral mistake. While courts are less likely to grant relief in cases of unilateral mistake than they are in cases of mutual mistake, “[t]here is practically universal agreement that, if the material mistake of one party was caused by the other, either purposely or innocently, or was known by the other or was of such character and accompanied by such circumstances that the other had reason to know of it, the mistaken party has the power to avoid the contract.”⁹⁷

While NBLLN is a well-recognized and pervasive behavioral bias, it is unclear whether a court would use the doctrine to set aside an NBLLN

95. Schwartz, *supra* note 77, at 1410.

96. *See supra* Subpart II.B.2.

97. 7 JOSEPH M. PERILLO, CORBIN ON CONTRACTS § 28.41 (rev. ed. 2002).

consumer's acceptance of a data privacy agreement. The discussion in the prior Part makes clear that the NBLLN consumer does not fully appreciate what she is giving away.⁹⁸ As previously discussed, it is as though she *thought* she was only trading away one apple, when in fact she was giving up a whole bushel of them.⁹⁹ This can be analogized to the classic case of *Hume v. United States*, in which the U.S. government entered into a contract to purchase shucks from Hume for 60 cents per pound.¹⁰⁰ According to the Supreme Court, this was a clerical error, and the intended price was sixty cents per hundred weight (i.e., per hundred pounds, or 0.6 cents per pound).¹⁰¹ The Court held that if Hume knew or should have known about the error, the contract should be set aside.¹⁰²

While a cognitive error is distinct from a clerical one, it is not obvious that they should be treated differently by the law. Moreover, given the overwhelming evidence of the NBLLN bias in the population, ISPs should be aware that consumers are unable to properly estimate the degree of privacy loss that they will experience, just as Hume should have known about the clerical error.

Despite the intuitive appeal of this doctrinal approach, we do not believe that the doctrine of unilateral mistake is an appropriate solution to the problem. In this context, it suffers from the same limitations as the doctrine of unconscionability.¹⁰³ Courts are poorly placed to evaluate the sufficiency and clarity of privacy disclosures, particularly in the context of the *other* information already available about an individual which must be aggregated to properly evaluate the mistake. Moreover, even if they were, nullifying a contract *ex post*, once the information has been transmitted and the privacy lost, is far from an ideal solution.

98. *See supra* Part II.

99. *See supra* Part III.A.

100. *Hume v. United States*, 132 U.S. 406, 407 (1889).

101. *Id.*

102. *Id.* at 414–15 (“If the claimant knew that a clerical error had been committed, of which the agents of the government were ignorant, and deliberately intended to take advantage of the error to obtain the execution of a contract for the payment of so grossly unconscionable a price, or if the facts were such that he must be held to have known that their action, if understandingly taken, would be in palpable dereliction of their duty to their principal, and, notwithstanding, sought to profit by it, the character of the fraud, so far as the claimant is concerned, is not changed by the fact that such action was the result of the negligence or mistake of the government’s agents, untainted by moral turpitude on their part.”).

103. Indeed, as the quote in footnote 102 makes clear, the court in *Hume* interpreted the doctrine of unilateral mistake as a manifestation of the unconscionability doctrine. *Id.* at 414.; *see also* 7 JOSEPH M. PERILLO, *supra* note 97 (“*Hume* demonstrates that relief for unilateral mistake descends from the doctrine of unconscionability.”).

C. REGULATORY APPROACHES

A much more straightforward approach is to regulate consumer privacy directly.¹⁰⁴ As previously discussed,¹⁰⁵ ISPs currently face far less regulation than companies such as Alphabet and Facebook, despite the fact that ISPs are positioned to collect far *more* consumer data. This is problematic. The FCC implicitly recognized this, and crafted the now-disapproved FCC Privacy Order. While a loss for consumer privacy, this action presents an opportunity to craft a more effective regulatory structure. In addition to relying on default rules, this structure should address consumers' NBLLN.

Because Congress disapproved the FCC Privacy Order through the Congressional Review Act, the FCC and other federal agencies are now barred from putting forth similar regulations.¹⁰⁶ Regulations addressing consumers' NBLLN in the internet privacy context are likely to fall into this category. As a result, such a regulation would need independent Congressional authority.

Work by Bar-Gill and Ben-Shahar has shown that default rules are largely ineffective in the context of consumer contracts.¹⁰⁷ This is particularly true in the privacy context.¹⁰⁸ Bar-Gill and Ben-Shahar give several explanations for this. Crucially, they note that consumers may simply lack the information that they need in order to make adequate opt-out decisions.¹⁰⁹ Strandburg has also argued that the personal data market works differently from standard markets.¹¹⁰

In contrast to Bar-Gill and Ben-Shahar, the problem that we call attention to operates even when individuals have all the necessary information, and they are fully aware of their preferences. In contrast to theirs, our model is built on the assumption that individuals suffer from a specific cognitive bias—NBLLN. As such, Bar-Gill and Ben-Shahar invoke a number of different deviations from rationality, while we are relying entirely on one specific one (that individuals suffer from NBLLN). On the one hand, this means that our analysis rests on the validity of that assumption. On the other hand, the parsimony of our assumptions

104. Strandburg, *supra* note 15, at 165–72 (arguing that neither notice and choice nor a more robust consent regime can overcome the basic problems of behavioral advertising business models).

105. *See supra* Part I.B.

106. *See supra* Part I.A.

107. *See* Oren Bar-Gill & Omri Ben-Shahar, *Optimal Defaults in Consumer Markets*, 45 J. LEGAL STUD. S137, S138–39 (2016) (noting that that consumers often ignore their own preferences and may not always understand default provisions).

108. *See id.*

109. *Id.*

110. *See* Strandburg, *supra* note 15, at 130–52 (“[U]nlike the payment of money in an ordinary retail transaction or the disutility imposed by broadcast or contextual advertising, data collection does not occur at a ‘point of purchase.’”).

means that we are more insulated from the standard critique of “fishing” for biases.¹¹¹ While we do not object to their assumptions, we do not rely on them in our framework.

Our framework shows that changing the default is not, on its own, enough. The problem is neither that consumers have limited attention (and are therefore not paying sufficient attention) nor that the default rules are “sticky” (perhaps because the default is interpreted as a suggestion, or because consumers simply do not understand the contract). Nor is it even necessarily that consumers do not know what information is already out there.¹¹² Rather, the problem is consumers’ inability to accurately estimate the incremental value of their information. Any direct regulation that addresses this issue should include provisions that reduce their NBLLN bias.

To see how to do this, we can return to the distinction we drew between transparency duties and consent duties in the FCC Privacy Order. In light of consumers’ NBLLN, any duties related to transparency must be designed in a way that heightens their ability to provide truly informed consent to an opt-in or opt-out.¹¹³ An NBLLN-robust disclosure is one that would allow such a consumer to understand what the information actually means in the context of all the *other* information being collected.

We will now sketch out the key attributes of an NBLLN-robust privacy disclosure. We do so in three sub-parts. The first—that the disclosure be clear and simple, and that it use examples that are easy for the reader to understand—is not unique to the NBLLN context. Rather, it follows existing “best practices” across the consumer contracting spectrum, and can be understood as a baseline upon which one must build the NBLLN-robust disclosure. It is, in other words, necessary but not sufficient. In contrast, our proposals in the next two sub-parts—how pieces of information combine and how classes of information accumulate—are tailored to the NBLLN context. Finally, the last sub-part brings these elements together, and discusses the limitations of existing proposals.

111. See discussion *supra* Part II.B.2.

112. See Strandburg, *supra* note 15, at 167 (explaining the problem of consumer uncertainty and stating that “[s]ince the market fails because of the impenetrability of data practices and the interconnectedness of information, the goal should be to do two things: incentivize data practices that are not impenetrable and disentangle the collection of data associated with different online activities,” and that “data practices should be such that consumers have an intuitive sense of what is going on with their data”).

113. See generally Jolls & Sunstein, *supra* note 64.

1. *Starting Point: Clear and Simple Disclosures*

The first step in achieving an NBLLN-robust privacy disclosure is describing the informativeness of the personal data collected from consumers in a way that is clear and easy to interpret. In other words, to address the accumulation problem that our model describes, one must first address the broader problem of disclosures that are written in terms that not even a fully rational individual could understand. This suggestion does not arise from our model but rather forms a backdrop for the remainder of this discussion. The notifications mandated in the FCC Privacy Order did not achieve this. While they would have required ISPs to disclose the data being collected from consumers, they would not, on their own, have been enough to solve the underlying information-processing problem in the face of NBLLN.

For notifications to be useful, they must address not only what information is collected but also the significance of such information. For consumers, it is not the same to read, “We will collect geo-location information” as it is to read, “We will collect geo-location information; this will tell us where you are accessing the internet from.” Moreover, even a non-NBLLN consumer is likely to benefit from clear and simple disclosure.

To make these disclosures effective, examples could be added about how the information collected can be used to learn about the consumer. While examples are useful to both NBLLN consumers as well as to non-NBLLN consumers, they may be particularly useful to consumers with NBLLN, since the main problem these consumers have is one of processing information correctly—they do not know how to aggregate the information available.¹¹⁴ Because examples are essentially “pre-processed” information, they can short-circuit this bias. For example, it is not the same for consumers to read, “We will collect geo-location information” as it is to read, “We will collect geo-location information; this will tell us where in the world you are when the app is running.” Like the use of clear and simple disclosures, the suggestion of using examples does not directly follow from our model. Rather, it forms part of a “baseline” upon which we build in the next two sub-parts.

2. *How Information Can Be Combined*

What is specific to our model is the way in which the data collected can be aggregated with *other* data to learn about the consumer. This is

114. See Strandburg, *supra* note 15, at 98 (explaining that “imperfect consumer information about the potential harms of data collection, company data practices, and means to mitigate data collection combine with the properties of information aggregation and with common behavioral economics concerns to undercut the market’s responsiveness to consumer preferences”).

where the problems specific to NBLLN begin to manifest themselves. Take the aforementioned example of geo-location. For consumers, it is not the same to read “we will collect geo-location information” as it is to read “we will collect geo-location information; this can be combined with geo-location information from your other devices and publicly available information and will tell us when you are at home, at work, or at a store.” The first does a good job at conveying the information clearly and reducing the information asymmetry between the company and the consumer about collection practices, but the second also explains the meaning of how different signals combine. With the first version, the consumer would know that her geo-location is being tracked, but she might not realize how easy it is to aggregate this with information about her home and work address to know how much time she spends at the office.

Part of why it is so difficult to calculate the value of one’s information is that, oftentimes, the value of information is not linear. Information often has synergies, and the value of a package of information is more than the sum of its parts. The nine digits of one’s social security number, for example, are much more valuable than nine times the value of each digit.¹¹⁵ Therefore, the value (or harm) of a particular piece of information can be very different depending on the exact combination of the *other* pieces of information that are already known.¹¹⁶

That being said, disclosure of what is *already known* is not enough on its own. While such a disclosure may help ameliorate consumers’ general confusion, it does little to mitigate their information aggregation problem. In other words, even if an NBLLN consumer was fully aware of everything that Poodle knew about her, she would *still* make the cognitive mistakes that cause her to undervalue her information.¹¹⁷

Instead, an effective disclosure regulation must help consumers understand how the *next* package of data will be used. For example, using geo-location once more, instead of “we will collect geo-location

115. For a similar observation, see Strandburg, *supra* note 15, at 130–152 (“[I]t is nearly impossible for a consumer to estimate the increment of expected harm associated with a given instance of data collection.”).

116. As others have pointed out, in addition to the problem that we focus on, consumers under the current regulatory regime have no way of knowing what information is already known about them. See Strandburg, *supra* note 15; see also Alessandro Acquisti & Ralph Gross, *Predicting Social Security Numbers from Public Data*, 106 PROC. NAT’L ACAD. SCI. U.S. 10975 (2009); Paul Ohm, *Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization*, 57 UCLA L. REV. 1701 (2010); Arvind Narayanan & Vitaly Shmatikov, *De-anonymizing Social Networks*, 2009 IEEE SYMP. ON SECURITY & PRIVACY 173 (2009); Arvind Narayanan & Vitaly Shmatikov, *Robust De-anonymization of Large Sparse Datasets*, 2008 IEEE SYMP. SECURITY & PRIVACY 111 (2008).

117. In fact, as noted above, our argument applies even when the NBLLN consumer is fully aware of what Poodle knows about her prior to the disclosure of interest. See *supra* Section II.C.

information,” consumers could read, “we will collect geo-location information; this can be combined with geo-location information that we collect from others, and can allow us to learn who you are spending time with, how much time, and where.” Just as the last example showed how to help a consumer aggregate different pieces of information that she already knew the company had, this disclosure helps her understand how the package of information aggregates with other information that the company will have—which might be more challenging for her if she is not paying attention. With the first disclosure, the consumer learns she is being tracked and, by extension, would know that her friends’ geo-location is being tracked as well, but she might overlook how these two pieces together can reveal much more meaningful information about her. These altered disclosures directly target the NBLLN bias identified in our model.

3. *Classes of Information*

Finally, consumers who are subject to NBLLN are also likely to misunderstand the interplay between different *classes* of information, which goes a step beyond understanding how to aggregate different packages of the same type of information. This is problematic: for example, the value of knowing exactly where the consumer is depends on whether the observer also knows what the consumer is *doing* when she is in each location.

Due to NBLLN, consumers make systematic errors when it comes to understanding how the different types of data fit together, and how clear of an image they can create. It follows from our model, then, that consumers will benefit from specific disclosure about how the information fits together. One example of this is what computer scientists call sensor fusion, which describes how different types of data collected from wearables and smart home devices aggregates to form new types of information that are hard to predict for a non-expert.¹¹⁸

For example, consider how the prior example of geo-location could combine with other types of information. Instead of adding a statement in the privacy notices, complying with the clarity requirement, stating that, “we will collect data on geo-location, web browsing history, app

118. See Scott R. Peppet, *Regulating the Internet of Things: First Steps Toward Managing Discrimination, Privacy, Security, and Consent*, 93 TEX. L. REV. 85, 120 (2014) (explaining that “[t]he technical problem created by the internet of [t]hings is that sensor data tend to combine in unexpected ways, giving rise to powerful inferences from seemingly innocuous data sources”); Andrew Raji et al., *Privacy Risks Emerging from the Adoption of Innocuous Wearable Sensors in the Mobile Environment*, CHI 2011 11, 11 (2011) (explaining that “seemingly innocuous data shared for one purpose can be used to infer private activities and behaviors that the individual did not intend to share.”).

usage history, and content of communications,” a consumer could read, “we will collect data on geo-location, web browsing history, app usage history, and content of communications. For example, this will allow us to learn where you are every time you open the app, and what websites you looked at before and after you used the app. It will also allow us to record the details of your app usage habits, when you communicate with other individuals through the app, and the full text of any conversations you have using the app platform.” This might sound like new information, but it is not. It is simply manifesting how the different types of information can be aggregated. This modification in the privacy notice would help consumers overcome the NBLLN and better understand the implications of a decision to disclose.

4. *Impact on Existing Notices Literature*

This discussion also points to an important implication in the extensive literature on the effectiveness of notices. While many scholars have called for more notices to consumers as a way to increase transparency,¹¹⁹ another stream of literature has suggested that notices do not effectively increase consumer awareness.¹²⁰ Indeed, empirical evidence has shown that simplifying disclosures has no effect on consumer awareness, suggesting that complexity in language is not the main driver.¹²¹ Moreover, other empirical work suggests that privacy language in itself is irrelevant, which in turn suggests that consumers do not react to different kinds of language.¹²²

119. See Ryan Calo, *Against Notice Skepticism in Privacy (and Elsewhere)*, 87 NOTRE DAME L. REV. 1027, 1047–59 (2011) (proposing visceral notices for privacy); see also Paula J. Dalley, *The Use and Misuse of Disclosure as a Regulatory System*, 34 FLA. ST. U. 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 COLUM. L. REV. 1701 (1999) (explaining the provision of notices as a common method for regulation in medicine).

120. See Solon Barocas & Helen Nissenbaum, *On Notice: The Trouble with Notice and Consent*, PROC. ENGAGING DATA F. (2009); Kirsten Martin, *Do Privacy Notices Matter? Comparing the Impact of Violating Formal Privacy Notices and Informal Privacy Norms on Consumer Trust Online*, 45 J. LEGAL STUD. S191 (2016) (using a vignette study to show that formal privacy notices actually reduce consumer trust on a website); see also Aleecia M. McDonald & Lorrie Faith Cranor, *The Cost of Reading Privacy Policies*, 4 I/S: J. L. & POL'Y FOR INFO. SOC'Y 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 BAYLOR L. REV. 139 (2006) (explaining the limits of a disclosure-based policy generally and suggesting direct conduct regulation through the example of securities).

121. See Omri Ben-Shahar & Adam Chilton, *Simplification of Privacy Disclosures: An Experimental Test*, 45 J. LEGAL STUD. S41, S65–66 (2016) (finding that best-practice simplification techniques have little or no effect on respondents' comprehension of disclosures).

122. See Lior Jacob Strahilevitz & Matthew B. Kugler, *Is Privacy Policy Language Irrelevant to Consumers?*, 45 J. LEGAL STUD. S69, S92–93 (2016) (testing language in privacy policies).

Our contribution provides a unified framework for understanding these results. Thus, the disagreement between proponents and detractors of notices as a means to help consumers could be due to the fact that effectiveness depends on the kind of notification and, more specifically, whether this notification effectively targets consumer bias.

CONCLUSION

Like most people in most situations, online consumers are not perfectly rational regarding their privacy choices. One central deviation from rationality is Non-Belief in the Law of Large Numbers. This leads consumers to make suboptimal choices in decisions that involve aggregating different pieces of information about them. In short, people are bad at estimating how these pieces of information combine. They suffer from a form of information overload, and end up with biased beliefs.

We discuss this behavioral fact in the context of ISP tracking and the disapproved FCC Privacy Order and demonstrate the extent to which consumers are unable to accurately estimate how much ISPs can learn about them based on their data. This fact provides both a foundation for regulatory intervention and suggestions for the form that these interventions should take.

This fact is relevant for any policy that wishes to address consumer privacy in a behavioral world. While it is particularly important in the ISP context, the implications of NBLLN for consumer privacy are relevant for all companies working with behavioral profiling. Forbidding the practice or forcing an opt-in consent will miss the mark, just as simply applying contract law principles is unlikely to be effective. Instead, a better way forward is through direct privacy regulations that enhance understanding. This approach would increase consumer welfare while maintaining profitable and legitimate business strategies.
