Separation, Pooling, and Predictive Privacy Harms From Big Data: Confusing Benefits for Costs

James C. Cooper*

DRAFT

October 7, 2015

Abstract: Privacy is concealing certain details about oneself from the world. In economic jargon, it means to "pool" with others. Economists, however, tend to prefer separating to pooling equilibria, as the former tend to lead to more efficient outcomes. This tension between privacy and market efficiency—between pooling and separating equilibria—is on full display in the burgeoning privacy law scholarship surrounding big data to make predictions about us. The privacy concerns raised in the big data context in large part have centered on so-called "predictive privacy harms," which arise as big data allows firms to make granular distinctions based on predictive algorithms and to tailor offers to customers, employees, and borrowers accordingly. The increasing use of algorithmic predictions based on big data have led some to call for limits on their use. Although informational asymmetries justified in the name of privacy are likely to give rise to costs, we cannot ignore that privacy itself has intrinsic value; most people are willing to pay to avoid unwanted surveillance, and privacy gives rise to broader social benefits. Indeed, information collection and concomitant algorithmic predictions designed merely to transfer surplus (e.g., "sucker lists") are wholl dissipative. Thus, merely to say that privacy retards information flows is insufficient to blunt calls for restrictions on the use of big data to make predictions about us. Nonetheless, just as policy should discourage dissipative investments in information revelation, it is equally crucial that policy discourages dissipative privacy—strategic concealment of facts relevant to a transaction in hopes of getting a better deal. This paper brings to bear insights from the economics of contracts and torts to develop a positive framework that helps identify dissipative and productive privacy, and that aids in identifying factors that militate for and against regulating big data. Application of the model to the use of big data to effect price discrimination and separation in labor and credit markets suggest—contrary to much of the scholarship—that the poor stand to gain from big data-driven separation. The extant empirical literature tend to support this prediction.

* Director of Research and Policy, Law & Economics Center, and Lecturer in Law, George Mason University School of Law. I thank Eric Claeys, Todd Zywicki, and participants at the George Mason University School of Law Levy Workshop for helpful comments. I also thank John Magruder, who provided outstanding research assistance.

INTRODUCTION

Privacy is about retarding information flows. It's about being "let alone,"[1] which boils down to being able to conceal certain details about oneself from the world. In economic jargon, it means to "pool". Economists, however, tend to prefer separating to pooling equilibria, as the former supply the market with more information and lead to concomitantly more efficient decisions. Markets create more surplus when information asymmetries between contracting parties are reduced: insurers can distinguish good from bad drivers, employers can discern productive from lazy workers; and lenders can distinguish good from bad credit risks. When parties conceal facts from one another, or cannot verify information that is voluntarily revealed, markets fail to reach their potential. They force "good" types to subsidize "bad types," and in the process attract too many bad types into a market (adverse selection). Private information also leads parties to spend resources to signal their true type or to establish mechanisms to tell good from bad. Because bad types do not shoulder their full weight, moreover, informational asymmetries can increase incentives to shirk (moral hazard). Indeed, several Nobel Prizes have been garnered by studying problems that arise from informational asymmetries in markets.[2]

This tension between privacy and market efficiency—between pooling and separation—is on full display in the burgeoning privacy law scholarship surrounding "big data," the use of the ever-growing data stream from online tracking and the so-called Internet of Things (IOT) to make predictions about us. The privacy concerns raised in the big data context in large part have shifted away from the more traditional domains of the unwanted collection, and concomitant risk of revelation of personal information like Social Security or bank account numbers. Rather, scholars and policy makers have begun to focus their attention on so-called "predictive privacy harms," which arise as big data allows firms to make granular distinctions based on predictive algorithms and tailor offers to customers, employees, and borrowers accordingly.[3] The sorting itself—and the concomitant diverse treatment— is the privacy harm. What tends to

---

[1] Samuel Warren & Louis Brandeis, *The Right to Privacy,* 4 HARV. L. REV. 193, 193 (1890).
[2] *E.g.*, George Akerlof; Joseph Stiglitz; Michael Spence.
[3] *See, e.g.*, Scott Peppet, *Unraveling Privacy: The Personal Prospectus & The Threat of a Full Disclosure Future*, 105 NW. U. L. REV. 1153 (2011); Ira S. Rubinstein, *Big Data: The End of Privacy or a New Beginning?*, 3 INT'L DATA PRIVACY L. 74 (2013); Kate Crawford & Jason Shultz, *Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms*, 55 B.C. L. REV. 93, 101 (2014); Edith Ramirez, Chairwoman, Federal Trade Commission, Opening Remarks at Big Data: A tool for Inclusion or Exclusion Workshop 4 (Sept. 15, 2014); Julie Brill, Commissioner, Federal Trade Commission, Big Data and Consumer Trust: Progress and Continuing Challenges 6, Remarks Before the International Conference of Data Protection and Privacy Commissioners (Oct. 15, 2014).

emerge is increasing calls to limit firms' ability to collect and use these data to make better inferences about those with whom they deal. It is not really an exaggeration to say that this line of scholarship views separation as a harm and pooling as the remedy.

At first glance, this policy prescription stands what most economists are taught in graduate school on its head—as a general matter, policies that force pooling are inefficient because they squander valuable information.[4] By the same token, we cannot ignore that privacy itself has value; most people would be willing to pay to avoid unwanted surveillance. Thus merely to say that privacy retards information flows is insufficient to condemn privacy-enhancing regulation. Privacy gains from pooling probably will outweigh efficiency gains from separation in many circumstances. Nonetheless, three important considerations are worthy of exploration.

First, one must be careful to distinguish between privacy's intrinsic and strategic values. Providing more of the former is welfare enhancing, while providing more of the latter is purely dissipative. By intrinsic value, I refer to the direct utility one derives from not being observed without consent—the value of limiting to oneself or a close circle of friends and family knowledge of certain personal facts. The strategic value of privacy, on the other hand, is the value that accrues to a party from obfuscating facts relevant to a transaction in hopes of getting a better deal. The prospective employee who hides the fact of her drug addiction is more likely to get the job; the prospective borrower who conceals his plans to quit his job is more likely to get a lower interest rate. It's important to emphasize that strategic privacy does not merely transfer value from employer to employee or from lender to borrower; it comes with real costs. Good types subsidize bad types by suffering lower wages and higher interest rates than they otherwise would, and overall welfare is reduced due to inefficient allocation of resources and wasteful expenditures on screening and signaling. Put differently, when privacy serves purely strategic purposes, losses to bad types due to big data-driven sorting should never be counted as a harm because they are merely artifacts of a net social benefit due to a reduction in adverse selection; without these losses, the net gains to society cannot materialize.

To the extent that a person's type is endogenous, insulating bad types from the impact of their decisions also dulls incentives to invest in becoming a good type. A host of empirical studies suggest that these adverse selection and moral hazard problems plague employment, insurance, and credit markets, and that the costs may be particularly acute those at the bottom of

---

[4] There are exceptions. When separating types leads only to distributional rather than productive gains, resources spent on separation—either signally or screening—are socially wasteful. *See* notes 80, *infra*, and accompanying text.

the economic rung who are relatively better risks than the rest of their cohort.[5]

In some cases it will be easy to distinguish between the strategic and intrinsic privacy values. The primary value from keeping traits like impulsiveness, poor work ethic, or lack of driving skill secret is likely to be strategic. In other cases, it may be difficult to disentangle the inherent and strategic dimensions of privacy. For example, most people probably derive utility from keeping health conditions private. But health privacy also has a strategic dimension; keeping the fact of diabetes or bi-polar disorder secret likely will lead to advantages in the labor and insurance markets.

A second consideration is that even when we can isolate instances where intrinsic, rather than strategic privacy is at stake, the current state of knowledge leaves policy makers ill-equipped to make rational tradeoffs. The thin extant literature on valuation of privacy provides little guidance on what the costs of policies that retard big data are likely to be.[6] What's more, the inherent valuation of privacy will vary across individuals and within individuals across contexts. A forty-six year old academic may be less willing to share the details of his weekend on Instagram than a twenty-year old college student who regularly tweets and posts the details of his life. Further, someone may feel strongly about keeping her health information private but couldn't care less about the public revelation of her Netflix queue. Others may have reverse preferences., The larger the variance in tastes for privacy, however, the more costly is a uniform rule. Absent a clearer understanding of consumers' intrinsic value of keeping certain information private, policy makers tackling big data should err on the side of caution. For example, it's unclear that the thousands of pregnant women who received discounts on diapers, cribs, and prenatal vitamins from Target would be willing to forego lower prices in return for not being classified as pregnant by a faceless algorithm.[7] Some probably would,[8] but there is a distribution, and we don't know it. Nonetheless, this episode is held up as exhibit A for predictive privacy harms from big data run amuck.[9]

---

[5] *See* parts III.A-B, *infra*, and accompanying text.

[6] *See* notes 96-108, infra and accompany text.

[7] *See* Kashmin Hill, *How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did*, Forbes (Feb. 16, 2012), http://www.forbes.com/sites/kashmirhill/2012/02/16/how-target-figured-out-a-teen-girl-was-pregnant-before-her-father-did/; Jordan Ellenberg, *What's Even Creepier than Target Guessing that Your Pregnant?,* Slate (June 9, 2014), http://www.slate.com/blogs/how_not_to_be_wrong/2014/06/09/big_data_what_s_even_creepier_than_target_guessing_that_you_re_pregnant.html.

[8] *See, e.g.*, Sarah Grey, *One Woman's Attempt to Hide Her Pregnancy from Big Data – it's More Difficult than You'd Expect*, Salon.com (Apr. 28, 2014), *at* http://www.salon.com/2014/04/28/one_womans_attempt_to_hide_her_pregnancy_from_big_data/.

[9] *See, e.g.*, Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995 (2014).

Finally, the distributional impacts of big data merit a closer examination. A common theme in the privacy literature is that big data disproportionately will harm the poor. Most of these worries, however, tend to fade when confronted with economic theory and empirical evidence: more granular predictions within sub-prime distributions can allow those relatively better credit risks to be identified and offered better terms; big data-driven predictions about employability can obviate the need to invest in educational signaling, which the poor often cannot afford. What's more, the price discrimination that many fear big data will make possible has the real potential to lower, not raise prices for the poor. It also has the potential to enhance competition for everyone if it allows firms to identify and poach—with lower prices—consumers of competing brands.

This paper does not attempt to argue that restrictions on big data necessarily fail a cost-benefit analysis given its promise to transform society—for example, through developing better medical treatments, reducing crime, or leading to more efficient farming techniques.[10] Instead, this paper is trained narrowly on the impact of big data restrictions on the relationship between the observed and the observer—or more precisely, the one who makes or acts on predictions from observed data. The overarching goal of this paper is to provide a positive framework for thinking about big data regulation that helps identify dissipative and productive privacy, and that injects a more fulsome understanding of the benefits that accrue from reducing asymmetric information into the debate. Specifically, I draw on the economic theory of contracts and torts to develop a simple model of optimal regulation when privacy harms are suffered heterogeneously. Along the way, I also attempt to address some of the concerns that big data is likely to have a disproportionate impact on the economically disadvantaged by bringing economic theory and empirical evidence into a debate that until now has been driven almost solely by anecdote and hypothetical.

This remainder of this paper is organized as follows. Part I describes big data and examines some of the privacy scholarship calling for its regulation based on so-called "predictive privacy harms." Part II examines the impact of information asymmetries on markets, including the qualities of separating and pooling equilibria. Drawing on the distinction between strategic and intrinsic values of privacy, Part III sets out a framework for analyzing privacy harms from big data. Part IV addresses claims that big data is likely to a disproportionally negative impact on the poor, and shows that there is reason to believe that the opposite is true. Part V concludes.

I. THE PROBLEM

---

[10] SeeVIKTOR MAYER-SCHONEBERGER & KENNETH CUKIER, BIG DATA (2013).

Much has been written on the topic, so I will only briefly describe big data to lay the groundwork for the remainder of the paper. Big data is a general catchall term for the analysis of enormous datasets – sets that may even satisfy the condition that "*N = all*"[11] – to tease out correlations and relationships that could not be seen with small data sets.[12] The rise of big data is made possible by the confluence of two factors: increasing digitization of our world, and increasing computing power. Words, sound, and video increasingly exist as zeros and ones, easy for computers to manipulate and analyze.[13] Computing storage and processing speed has grown in tandem with this increase in data, so that now we are able to analyze large data sets.

Google FluTrends, a Google-invented algorithm that predicts flu outbreaks based on Google search terms, is often held out as the quintessential example of big data.[14] Other notable examples include Oren Etzioni's Fare Cast flight price prediction web site,[15] Amazon's book recommendation algorithm, Google Translate,[16] and Netflix's movie recommendation algorithm.[17] Credit card companies also use big data methods to detect fraud by examining anomalies in purchasing patterns.[18] As will be discussed in more detail below, big data is also making inroads into finance and employment markets, where companies are using a variety of traditional and non-traditional data sources to make predictions about creditworthiness and job-suitability.[19]

---

[11] *Id.* at 26; *see also Id*. at 6 (big data refers to "things that one can do at a large scale that cannot be done at a smaller one, to extract new insights or new forms of value . . .").
[12] *Id*. at 12 (big data "is about applying math to huge quantities of data in order to infer probabilities").
[13] STEPHEN DOYLE, ESSENTIAL ICT A LEVEL: AS STUDENT BOOK FOR AQA 131 (2008).
[14] *See* GOOGLE FLUTRENDS, http://google.org/flutrends/us/#US; *But see* Paul Ohm, *Response: The Underwhelming Benefits of big data*, 161 U. PA. L. REV. 339 (2013), http://www.pennlawreview.com/online/161-U-Pa-L-Rev-Online-339.pdf.
[15] Farecast was purchased by Microsoft and integrated into Bing, but recently shut down. *See Farwell Farecast: Microsoft Kills Airfare Price Predictor, to the Dismay of its Creator*, Geekwire (Apr. 18, 2014), http://www.geekwire.com/2014/farewell-farecast-microsoft-kills-airfare-price-predictor-dismay-creator/.
[16] Tim Harford, *Big Data: Are We Making a Big Mistake?*, FINANCIAL TIMES (Mar. 28, 2014)("Google Translate is as close to theory-free, data-driven algorithmic black box as we have"), *at* http://www.ft.com/intl/cms/s/2/21a6e7d8-b479-11e3-a09a-00144feabdc0.html.
[17] *See* BIG DATA at 110.
[18] *See Crunching the Number*, THE ECONOMIST (May 19, 2012).
[19] *See* part IV.B., infra, and accompanying text.

A great deal has been written attempting to classify harms that stem from data breaches or the unwanted collection and use of personal data.[20] There is little trouble in classifying lost money from stolen credit card numbers or hacked bank accounts as harm. Similarly, even if identity theft does not result in direct financial losses, the time and hassle of reestablishing one's identity clearly is harmful. Then there are subjective harms, which include online tracking by ad networks, breach of health care information, or unwanted surveillance of intimate activities.[21] Because they are not objectively verifiable like monetary harms, and are likely suffered heterogeneously across populations and contexts, they are difficult to quantify. Nonetheless, they are harms in the traditional sense.

As big data has become the major focus of privacy discussions, the concept of harm has shifted. Although privacy scholars continue to raise traditional privacy concerns associated with unwanted collection and use of personal information, including the specter of easy re-identification and that large data reservoirs will make easy targets for hackers,[22] an increasingly popular target is classification made possible by big data analytics. In this manner, the focal point of big data privacy is the picture of oneself that emerges when a torrent of seemingly innocuous bits of data from the real and virtual worlds are run through predictive algorithms, and how this picture is used. This picture may be quite personal – like the transporter on the Enterprise, when the algorithm reassembles these tiny bits of data into a "person," it may reveal private aspects of one's life that many would not divulge publicly, like sexual orientation, drug use, or health status. Once the data have spoken, firms will be able to tailor offers based on your reconstructed person.

Some have expressed concern over the potential big data has to make discrimination easier. For example, bigots could hide behind impersonal algorithms pre-baked to exclude women or minorities. Others have expressed concern that even if not consciously used to discriminate, big data driven algorithms nonetheless might end up making classification along racial or gender lines.[23] Leaving aside the ability of big data to facilitate discrimination against protected classes, some privacy scholars also bemoan the use of big data-driven predictions by firms to customize prices, credit

---

[20] *See*, *e.g.*, Daniel Solove, *A Taxonomy of Privacy* 154 U. PA. L. REV. 477 (2006); Ryan Calo, *The Boundaries of Privacy Harms*, 86 IND. L.J. 1131 (2012).
[21] *See, e.g.*, Compliant, *In re Designerware, LLC*, Docket No. C4390 (F.T.C. April 15, 2013), *available at* https://www.ftc.gov/sites/default/files/documents/cases/2013/04/130415designerwarecmpt.pdf.
[22] *See, e.g.*, Ohm, *supra* note 12, at 341 (arguing that Google's use of search queries without permission violated privacy rights); Dennis D. Hirsch, *The Glass House Effect: big data, The New Oil, and The Power of Analogy*, 66 ME. L. REV. 373, 375 (2014).
[23] *See* Crawford & Shultz, *supra* note 3, at 101; Elizabeth Dwoskin, *How Social Biases Creeps into Web Technologies,* WALL STREET JOURNAL (Aug. 21, 2015).

offers, insurance rates, or employment opportunities.[24] They paint a dystopian future where economic opportunities – employment, prices, credit – are based on ubiquitous monitoring of all aspects of life. This use of big data is the focus of this paper.

A recent piece by Scott Peppet expressing concern over the "unexpected inferences about individual consumers" that may arise from big data is representative of this literature.[25] Importantly, the worry is not that the data or inferences will be inaccurate, but rather that big data will tell too much:

> Employers, insurers, lenders, and others may then make economically important decisions based on those inferences, without consumers or regulators having much understanding of that process. This could lead to new forms of illegal discrimination against those in protected classes such as race, age, or gender. More likely, it may create troublesome but hidden forms of economic discrimination based on Internet of Things data.[26]

Peppet allows that these sorts of big data-driven separating equilibria are likely to create efficiencies, but nonetheless cautions "from a legal or policy perspective, however, economic sorting is just not that simple" because "the public and its legislators tend to react strongly to forms of economic discrimination."[27]

Other scholars in this vein similarly have labeled instances in which big data is used to sort consumers into categories as "predictive privacy harms" or "classification harms."[28] Target has become the poster child of

---

[24] One notable exception is the work of Lior Strahilevitz, which sees the possibility for big data to decrease discrimination against protected classes. He reasons that to the extent that discrimination is motivated by economic, as opposed to insidious, reasons, more accurate information about a person's characteristics will reduce the use of protected status as a proxy. For example, if prison records are available, employers will stop using race as a proxy for the probability of past imprisonment, likely improving the prospects of applicants from races with disproportionally high imprisonment rates. *See* Lior Jacob Strahilevitz, *Privacy vs. Antidiscrimination*, 75 U. CHI. L. REV. 363, 376 (2008).

[25] *See* Scott Peppet, *Regulating the Internet of Things: First Steps Towards Managing Discrimination, Privacy, Security, and Consent*, 93 TEX. L. REV. 85 (2014).

[26] *Id.* at 28.

[27] Peppet, *supra* note 21, at 36.

[28] Kate Crawford & Jason Shultz, *Big Data and Due Process: Toward a Framework to Redress Predictive Privacy Harms*, 55 B.C. L. REV. 93, 101 (2014) (because big data predictions "create a model of possible personal information and associate it with an individual, . . . harms can result regardless of the model's accuracy"); Cynthia Dwork & Deirdre K. Mulligan, *It's Not Privacy and It's Not Fair*, 66 STAN. L. REV. 35, 36 (2013), http://www.stanfordlawreview.org/online/privacy-and-big-data/its-not-privacy-and-its-not-fair (noting "concerns with the classifications and

classification harms for using analytics to send coupons for maternity-related products, such as diapers, cribs, and pre-natal vitamins to customers whose shopping habits suggested a high likelihood of being pregnant.[29] The supposed harm was that Target was using bits of innocuous public data—shopping habits—to construct a prediction of personal data—pregnancy status—and then to use this categorization to target discounts.[30] Even when the categorization is not based on a prediction of personal data, the mere fact that categorization leads to winners and losers is sufficient cause for alarm to some.[31] In another related and widely cited article, Ryan Calo expresses concern that firms will use big data algorithms to detect those who exhibit behavioral biases and take advantage of them.[32] He argues that firms will use big data to charge consumers "as much as possible" and to manipulate them to buy products and services that they "[do] not need or need[] less of."[33] For example, Calo suggests that a company could use big data to send junk food offers to those who big data has determined suffer from a lack of will power.[34]

A common theme in much of this work is that big data classifications overwhelmingly benefit the rich at the expense of the poor. For example, it has been suggested that big data will be used to offer discounts to the rich on luxury goods, which are subsidized by high prices for the poor on staples, like bread and milk.[35] Crawford & Shultz lament the fact that rich and poor

---

segmentation produced by big data analysis"); Ira Rubinstein, *Big Data: The End of Privacy or a New Beginning*, 3 INT'L DATA PRIVACY L. 65 (2013).

[29] *See, e.g.*, Crawford & Shultz at 98-99; Tim Harford, *Big Data: Are We Making a Big Mistake?*, FIN. TIMES (Mar. 28, 2014).

[30] Although the outrage has been round, it is unclear whether the critics would prefer Target to provide prenatal vitamin discounts to everybody or nobody.

[31] *See* Joesph W. Jerome, *Buying and Selling Privacy: Big Data's Different Burdens and Benefits*, 66 STAN. L. REV.47, 51 (2013) ("In the end, the worry may not be so much about having information gathered about us, but rather being sorted into the wrong or disfavored bucket."); Omer Tene & Jules Polensky, *Judged by the Tin Man: Individual Rights in the Age of Big Data*, J. ON TELECOMM.& HIGH TECH. L. 351, 367 (2013) ("A better understanding of the effect of data analysis on fairness, discrimination, siloization and narrowcasting can expand the scope of privacy harms that are subject to legal protections."); Omer Tene & Jules Polensky, 11 *Big Data for All: Privacy and User Control in the Age of Analytics*, 11 NW. U, L. REV. 240 (2013). In earlier work, Peppet explains how consumer signaling may substitute for firm screening as consumers/employees/lenders increasingly will be able to provide credible information about their type with sensor data. Of course, because those with favorable data to report will want to report, companies naturally will infer that non-reporters are of the "bad" type. Peppet's concern is that the increasing ability to credible reveal one's type will reduce privacy by raising the price of non-revelation. *See* Peppet, *supra* note 3. More generally, Dwork and Mulligan lament the potential for big data to create "filter bubbles" that "create feedback loops reaffirms and narrowing individuals' worldviews." *See* Dwork & Mulligan, *supra* note 25, at 37.

[32] Calo, *supra* note 8.

[33] *Id.* at 33.

[34] *Id.* at 31.

[35] *See* Omer Tene, *Privacy: For the Rich or for the Poor*, CONCURRING OPINIONS (July 2012), http://concurringopinions.com/archives/2012/07/privacy-for-the-rich-or-for-the-poor.html

receive different credit offers online.[36] Further, Joseph Jerome concedes that big data will enhance market efficiency, but nonetheless warns "market efficiency favors the wealthy, established classes."[37] He adds "categorization and classification threaten to place a privacy squeeze on the middle class as well as the poor."[38]

Not surprisingly, these authors generally recommend a government response to the problems posed by big data.[39] Peppet, for example, suggests limiting consumers' ability to acquiesce to monitoring via IOT sensors to avoid the negative inference that a firm could draw about one's type from an unwillingness to be monitored.[40] Further, some authors have suggested "due process" rights in big data predictions, likening these determinations to government deprivations of liberty. For example, Crawford & Shultz propose the a sliding scale of due process requirements, depending on the type of "predictive privacy harm:" determinations involving health would receive the most protection; advertising would receive less scrutiny, whereas "mixed uses" involving both advertising and health information like the Target pregnancy debacle, would receive the same protection as health information.[41] This protection would include some form of notice over what data is going into the classification scheme and the ability to challenge the fairness of a big data classification before an impartial adjudicator.[42]

Concern over big data's potential to classify people is not just academic. The Chairwoman of the FTC, for example, has warned of what she call's "data determinism," which occurs when individuals are judge "because of inferences or correlations drawn by algorithms suggest that they may behave in ways that make them poor credit or insurance risks, unsuitable candidates for employment or admission to schools or other institutions, or unlikely to carry out certain functions."[43] Her colleague,

---

(discussing a paper by Laura Moy & Amanda Conley, *Paying the Wealthy for Being Wealthy: The Hidden Costs of Behavioral Marketing*).

[36] *See* Crawford & Shultz, *supra* note 3, at 101.

[37] Jerome, *supra note* 28, at 50.

[38] *Id.*

[39] *See, e.g.*, Peppet, *supra* note 21, at 58-62 (arguing for restrictions on "cross-context" use of data streams and analogizing them to FCRA, the 5th Amendment, and the Genetic Information Nondiscrimination Act); Peppet, *supra* note 2, at 48-49; Dwork & Mulligan, *supra* note 25, at 39 (suggesting the establishment of a metric "defining who must be treated similarly" that "creates a path for external stakeholders . . . to have greater influence over, and comfort with, the fairness of classifications.").

[40] Peppet*, supra* note 21, at 56.

[41] Crawford & Shultz, *supra* note 3, at 118.

[42] *Id.* at 126-28.

[43] *See* Edith Ramirez, Chairwoman, Federal Trade Commission, The Privacy Challenges of Big Data: A View from the Lifeguard's Chair 8, Keynote Address at the Technology Policy Institute Aspen Forum (Aug. 19, 2013). *See also* Ramirez, *supra* note 3 (warning of using big data to

Commissioner Julie Brill, similarly has expressed concern that "the same data that allows banks to reach the traditionally unbanked, financially vulnerable populations could just as easily be used to target them with advertisements for high-interest payday loans."[44] The recent FTC report on Data Brokers echoed these apprehensions over classification, such as if an insurance company used information suggesting risky behavior or diabetes to adjust premiums,[45] and it recommended Fair Credit Reporting Act (FCRA)-like legislation to cover data brokers.[46] Further, this spring, the White House floated a draft privacy bill that adopted a strong regulatory stance toward big data predictions.[47] For example, data analysis that has the potential to result in "adverse actions concerning multiple individuals," would require a disparate impact analysis, and "privacy review boards" would be tasked to consider "professional harm" as a cost to be weighed against benefits when determining whether a data practice passes muster.[48]

The extant literature gives lip service to the economic efficiencies that are likely to flow from big data's ability to make the world less opaque, but quickly dismisses them as secondary compared to predictive privacy harms.[49] Clearly, consumers value privacy and it may be that privacy concerns ultimately rule the day. Big data's potential to reduce information asymmetries, however, needs to be taken seriously before one can call for regulatory intervention. That task is taken up in the next part.

II. ASYMMETRIC INFORMATION: ADVERSE SELECTION, MORAL HAZARD, & BIG DATA

At the end of the day, those concerned with classification harms really are concerned with big data's potential to promote separating equilibria.[50] Such concerns, however, run contrary to the general proposition that separation is better than pooling. Because it lies at the heart of the matter, it is useful to explore these concepts in some detail.

---

segment along income or racial lines, and referring to this practice as "discrimination by algorithm" and "digital redlining").

[44] Brill, *supra* note 3.

[45] Federal Trade Commission, Data Brokers: A Call for Transparency and Accountability 48 (2014).

[46] *Id.* at 51-52. Further, Chairwoman Ramirez and Commissioner Brill also support requiring data brokers to assure that their data sources acquired the data through "notice and choice, including express affirmative consent for sensitive data." *Id.* at 52 n.91.

[47] *Administration Discussion Draft: Consumer Privacy Bill of Rights Act of 2015*, available at https://www.whitehouse.gov/sites/default/files/omb/legislative/letters/cpbr-act-of-2015-discussion-draft.pdf.

[48] *See id.* at Sec. 103.

[49] *See* Peppet, *supra* note 3; Jerome, *supra* note 28, at 51.

[50] Peppet, *supra* note 21, at 41.

A. *Separation, Pooling and Adverse Selection*

Heterogeneity is a fact of life. People differ over myriad dimensions that are not directly observable, such as intellect, work ethic, maturity, and impulsiveness. In a world of perfect information, contracts would reflect these differences: those least likely to default would have greater access to credit and pay lower interest rates; those least likely to suffer an accident would have higher insurance levels and pay lower premiums; and those with greater work ethics would get better jobs and earn higher wages. Problems arise, however, because these traits are private information and can be difficult to verify. As a result, such markets can be characterized by adverse selection, which occurs when a firm's offerings attract a disproportionate amount of "bad" types – e.g., risky borrowers, unproductive workers, bad drivers, those with unhealthy lifestyles, and the like.

Take the canonical example of *Hadley v. Baxendale*.[51] There are two types of millers—those with a spare shaft (good types), who will continue to operate when one breaks, and those without (bad types), who will be down until the broken shaft is repaired. *Ex ante*, the courier hired to take the broken shaft for repair has no way of identifying one miller type from the other, so in a pre-*Hadley* world he charges an average price based on expected damages in the event he breaches. Two-shafters would gain by identifying themselves, but in this example so few millers have only one shaft that the gap between the average price and the two-shaft price is too small to make it worthwhile. As recognized by Ayers & Gertner in their classic article, a rule allowing unforeseeable consequential damages in these circumstance will create incentives for the one-shafters to hide among—or *pool with*—the two-shafters.[52] This price is a bargain for one-shafters; the receive insurance from the courier at a price subsidized by two-shafters, who will claim below average damages in the event of breach. This type of cross-subsidization is the hallmark of adverse selection. Bad types are drawn into the market because they can free ride off of good types. Of course, this causes goo types to leave, resulting in a market characterized by lower output a greater proportion of bad types than would exist with full information. In the extreme, adverse selection can cause markets to unravel completely.[53]

[51] *Hadley v.* Baxendale, 156 Eng. Rep. 145 (Ex. Ch. 1854).
[52] Ian Ayers & Robert Gertner, *Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules*, 99 YALE L.J. 87 (1989).
[53] George Akerlof, *The Market for Lemons: Quality Uncertainty and the Market Mechanism*, 84 Q. J. ECON. 488 (1970).

Adverse selection can be found in a variety of markets in which one party is likely to have private information.[54] Employers, lenders, or insurers observe proxies for latent qualities – employers can read college transcripts and talk to past employers, lenders can verify employment and look at credit scores, auto insurers look at age, employment, and past driving experience. But even within a group that looks homogenous across a variety of observable traits, there are likely to be important latent differences that impact the value of the contractual relationship.[55] Two potential employees may look similar on paper, for example, but one views the job merely as a weigh-station while his spouse finishes medical school. Two potential borrowers may have similar incomes and credit scores, but one knows that she is in an unstable marriage and is planning to quit her job in two weeks for a speculative work-from-home opportunity. What's more, adjusting the price to reflect average risk can exacerbate adverse selection. For example, insurers must grapple with the fact that those who are most likely to make claims are precisely the consumers who are most willing to purchase the most insurance coverage at the highest rates.[56] In credit markets, lenders understand that higher interest rates will attract a disproportionate share of consumers who are more likely to default. And employers who offer lower wages risk attracting only the least productive workers.

Firms ideally would like to find a way to separate good from bad types and offer each a contract that reflects their true types. One strategy is to

---

[54] *See* Lawrence M. Ausbel *Adverse Selection in the Credit Card Market*, (1999) (credit card markets); Wendy Edelburg, *Risk-Based Pricing of Interest Rates for Consumer Loans,* J. OF MONETARY ECON. (2006) (consumer loan market); Liran Einav, Mark Jenkins & Jonathan Levin, *The Impact of Credit Scoring on Consumer Lending*, 44 RAND J. OF ECON. 249 (2013) (subprime auto loan market); Adams et al. (AER 2009) (subprime auto loan market); Bev Dahlby, *Testing for Asymmetric Information in Canadian Automobile Insurance* (1992) (auto insurance); Daniel Altman, David M. Cutler & Richard Zeckhauser, *Adverse Selection and Adverse Retention*, 88 AM.ECON. REV. 122 (1998) (health insurance); Amy Finkelstein & James Poterba, *Testing for Asymmetric Information Using 'Unsused Observables' in Insurance Markets: Evidence from the U.K. Annuity Market*, J. OF RISK & INS. (Dec. 2014); Dean Karlan & Jonathan Zinman, *Expanding Credit Access: Using Randomized Supply Decisions to Estimate the Impacts*, (2008), http://karlan.yale.edu/p/Karlan&Zinman_ExpandingCreditAccess_jan2008.pdf (South African subprime lender); Robery Puelz & Arthur Snow, *Evidence on Adverse Selection: Equilibrium Signaling and Cross-Subsidization n the Insurance Market*, 102 J. OF POL. ECON. 236 (1994). *But see* Pierre-Andre Chiappori & Bernard Salanie, *Testing for Asymmetric Information in Insurance Markets*, 108 J. OF POL. ECON. 56 (2000) (no evidence in French auto insurance market for first time drivers); James H. Cardon & Igal Hendel, *Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey*, 32 RAND J. OF ECON. 408 (2001) (health insurance).

[55] This is the point behind esurance's "Sorta like you isn't you" campaign. See https://www.esurance.com/quote1301?PromoID=GGNBB_VA_001&partner_cd=AdPos-1t1%7CGeo-9008162%7CAdID-79094026107%7C&ts=2.

[56] Part of the hidden information that leads to adverse selection can include intended *ex post* effort. For example, conditional on being insured, some may intend to consumer more health care than they otherwise would. *See* Karlan & Zinman, *supra* note 49; Einav, Jenkins & Levin, *supra* note 49; Finkelstein & Poterba, *supra* note 49.

screen potential customers by offering a set of contracts that will create incentive for types to reveal themselves. This is in essence what a good test does. Because only the best students will be able to answer a subset of the questions, it allows the professor to achieve separation and assign a distribution of grades that ostensibly reflects true mastery of the material. For such an equilibrium to be feasible, however, the firm must be able to offer a contract that is suboptimal (relative to the full information optimum) for the good type—but better than the option of exiting the market—to avoid adverse selection.[57] If the contract offered to good types is too favorable, it will attract both types, and prevent separation.[58] In this manner, parties cannot use price as the sole instrument to effect separation and instead must resort to rationing—e.g., down payments, deductibles, caps—to clear markets. As a result, good types bear too much risk or receive too little credit compared to the full information equilibrium.

The second way for separation to occur is for good types to reveal themselves with a signal. They clearly have an incentive to do so, but unfortunately they can't merely by declaring themselves good. This is "cheap talk"—a signal that's costless for either type to send, and hence conveys no credible information. Rather, for a signal to promote separation, it must be too costly for the bad type to send. In Spence's seminal job market signaling paper, for example, education can signal productivity only if high-productivity workers can acquire education sufficiently more cheaply than their low-productivity counterparts.[59]

Clearly, in markets characterized by adverse selection, bad types exert a negative externality on good types. When a separating equilibrium cannot be obtained because signaling or screening is too expensive relative to the gains, good types are forced to subsidize bad types in a pooling contract. Although separation is preferred, even when it can be obtained through screening or signaling, it comes at a price: good types bear too much risk,

---

[57] *See* Joseph E. Stiglitz & Andrew Weiss, *Credit Rationing with Imperfect Information*, 71 AM. ECON REV. 393 (1981).

[58] This is analogous to a test that has only easy questions, allowing the poor students to pool with the good.

[59] *See* Michael Spence, *Job Market Signaling*, 87 Q. J. ECON 355 (1973). There is also a host of empirical work that finds evidence that education serves as a signal in labor markets, which is consistent with asymmetric information in these markets. *See* Kelly Bedard, *Human Capital versus Signaling Models: University Access and High School Dropouts*, 109 J. Pol. Econ. 749 (2001); John H. Tyler, Richard J. Murnane & John B. Willett, *Estimating the Labor Market Signaling Value of the GED*, 115 Q. J. ECON. 431 (2000); David A. Jaeger and Marianna E. Page, *Degrees Matter: New Evidence on Sheepskin Effects in the Returns to Education*, 78 REV. OF ECON. & STAT. 733 (1996); Kevin Lang & David Kropp, *Human Capital versus Sorting: The Effects of Compulsory Attendance Laws*, 101 Q.J. ECON. 209 (1986); John G. Riley, *Testing the Educational Screening Hypothesis*, 87 J. POL. ECON. 227 (1979); Richard Layard & George Psacharopoulos, *The Screening Hypothesis and the Returns to Education*, 82 J. OF POL. ECON. 985 (1974).

receive too little credit, or receive low wages compared to a full-information equilibrium. Alternatively, they must invest in costly signaling. By increasing the price of good types participating the market, moreover, informational problems reduce overall output and welfare.

B. *Dynamic Considerations: Moral Hazard and Endogenous Types*

In addition to adverse selection, markets characterized by asymmetric information are often subject to moral hazard. Whereas adverse selection concerns hidden information about parties before they enter into a relationship, moral hazard concerns hidden actions—actions that impact the value of the relationship—that occur *after* the parties enter into a contract. If the party whose actions impact the value of the contract does not bear the full costs of these actions, there is a natural tendency to engage in suboptimal effort. For example, drivers have the ability to reduce the probability that they will get into an accident by choosing to drive more slowly, less often, and on less congested roads. When one is fully insured, however, they have less incentive to take these actions because they are costly. Borrowers have control over whether they will be able to repay their loan, for example, by restraining current spending and taking actions to ensure sufficient income flow. To the extent that a borrower can escape the full cost of default, they will take less take less care to avoid default, because these actions are costly. In this manner, moral hazard is the flip-side of adverse selection: adverse selection occurs when riskier individuals select into the market; moral hazard occurs when market participation increases incentives to engage in riskier actions.[60]

Parties take a variety of actions to ameliorate moral hazard. Insurers concerned about moral hazard, for example, require deductibles and have coverage limits. Lenders concerned with moral hazard limit loan amounts, and require down payments and other types of collateral. Both of these strategies involve rationing in the cause of creating incentives for consumers to take actions to avoid accidents or default. As is the case in the presence of adverse selection, this rationing is costly: consumers bear too much risk and have too little access to credit.

Not surprisingly, empirical evidence suggests moral hazard exists in lending and insurance markets. In a recent paper, for example, Karlan & Zinman find evidence of moral hazard in credit markets for poor South Africans.[61] Edleburg, moreover, finds evidence of moral hazard in U.S.

---

[60] *See* Chiappori & Salanie, *supra* note 49, at 60.
[61] Karlan & Zinman, *supra* note 49. *See also* Edelburg, *supra* note 49 (finding evidence of moral hazard in auto and credit card lending in the U.S.).

consumer lending markets. Several papers have also found evidence of moral hazard in insurance markets.[62] Moral hazard exists in other setting in which parties do not bear the full risk of their actions. For example, some empirical work suggests that consumers tend to take less care when using risky products if they are likely to be insured through products liability.[63] Finally, several studies document the so-called "Peltzman" effect, in which actors take less care when there are exogenous increases in safety.[64]

C. *Big Data*

It is not hard to see how big data could improve the performance of markets fraught with asymmetric information. To the extent that big data allows lenders, insurers, or employers to have a clearer picture of a person's type, it will reduce problems associated with adverse selection. In some cases, bad types will suffer as they are forced to pay the correct rate for their type, but in other cases they won't. Recall that achieving separation through screening or signaling requires a sacrifice from only the good type. Bad types enjoy the same terms they would in full information and do not have to invest in signaling. If big data can eliminate or reduce the need to employ these costly hurdles, then it is Pareto improving—good types gain and bad types are no worse off. In either case, greater separation in types enhances overall welfare by leading to more efficient matching.

For example, alternative credit scoring mechanisms use a variety of predictors—from social media posts, to payment of cell phone bills—to predict credit worthiness.[65] Individuals with stable networks of close friends and whose information on LinkedIn matches his application, or businesses

---

[62] *See* Puelz & Snow, *supra* note 49; Yingying Dong, *How Health Insurance Affects Health Care Demand-A Structural Analysis of Behavioral Moral Hazard and Adverse Selection*, 51 ECON. INQUIRY 1324 (2011). *But see* Chiappori & Salanie, *supra* note 49. Einav, Jenkins & Levin, *supra* note 49; Jonathan Klick and Thomas Stratmann, *Subsidizing Addiction: Do State Health Insurance Mandates Increase Alcohol Consumption?*, 35 J. OF LEGAL STUD. 175 (2006) (For example, in a series of papers Jon Klick and Thomas Strattman find evidence of moral hazard when state laws mandate coverage of diabetes and alcohol abuse treatment).

[63] Paul H. Rubin & Joanna M. Shepherd, *Tort Reform and Accidental Deaths*, 50 J. L. & ECON. 221 (2007). Similarly, Helland & Tabborock reveal evidence of moral hazard in general aviation, as they show that accidents fall and investments in safety by pilots increase as expected liability compensation falls. Eric A. Helland & Alexander Tabarrok, *Product Liability and Moral Hazard: Evidence from General Aviation*, 55 J. L. & ECON. 593 (2012).

[64] Sam Peltzman, *The Effects of Automobile Safety Regulation*, 83 J. POL. ECON. 677 (1975). See also John M. Yun, *Offsetting Behavior Effects of the Corporate Average Fuel Economy Standards*, 40 ECON. INQUIRY 260 (2002); Robert S. Chirinko & E.P Harper, Jr., *Buckle Up or Slow Down? New Estimates of Offsetting Behavior and Their Implications for Automobile Safety Regulation*, 12 J. POLICY ANALYSIS & MGM'T 270 (1993).

[65] Nate Cullerton, *Behavioral Credit Scoring*, 101 GEO. L.J. 808, 809 (2013).

with good reputations on social media are more likely to get loans.[66] These big data-driven sorting techniques are especially valuable to those consumers or businesses with little credit history.[67] Similarly, some employers are using big data predictions about potential employees to supplement, or even replace traditional hiring techniques for some jobs.[68] One firm has examined employee email, calendars, and HR record, and found a correlation between attendance of events and benefits coverage selection and the likelihood of an employee quitting within a year.[69] These techniques can reduce the substantial costs associated with worker churn. Further, Wal-Mart is reportedly using big data to predict who is likely to get promoted in an effort to limit the length of vacant jobs.[70]

What role might big data play in ameliorating problems associated with moral hazard? First, moral hazard can be tempered through separation. Bad types lose in separating equilibria because good types no longer subsidize them. However, in some cases types are not immutable characteristics, but instead a result of choices made under moral hazard. Recall the separation brought about by the court in *Hadley v. Baxendale*.[71] Although the one-shaft millers are worse off, society gains because pricing sends more accurate signals: they tell the courier to take more take more care with one-shaft millers, but they also for the one-shaft miller to bear the full cost of his decision to have only one shaft. To the extent that this decision was made because he was receiving free insurance from the courier, he now may finally buy that second shaft. Thus, although some portion of a person's risk profile may be exogenous, other components are endogenous; big data can impact the latter. If big data's unmasking of bad types forces them to pay prices that more closely reflect their true risk, they are likely to alter their behavior to the extent that it is feasible. For example, if alternative credit scoring limits pooling among sub-prime populations, it may incentivize the relatively worse credit risks to take steps to reduce the likelihood of missing a payment. Similarly, feeding big data predictions about healthy lifestyles from purchases and other trackable behaviors into insurance rates may reduce incentives to engage in unhealthy behaviors, such as smoking or sedentary lifestyles.

---

[66] *See* Stephanie Armour, *Borrowers Hit Social-Media Hurdles*, WALL STREET JOURNAL (Jan. 8, 2014).

[67] *Id*. *See also* notes 141-147, and accompany text, *infra*.

[68] Claire Cain Miller, *Can an Algorithm Hire Better than a Human?*, NEW YORK TIMES (June 25, 2015), *at* http://www.nytimes.com/2015/06/26/upshot/can-an-algorithm-hire-better-than-a-human.html?_r=0; Max Nisen, *MONEYBALL AT WORK: They've Discovered What Really Makes a Great Employee*, Business Insider (May 6, 2013), http://www.businessinsider.com/big-data-in-the-workplace-2013-5.

[69] Rachel Emma Silverman & Nikki Waller, *The Algorithm that Tells the Boss Who Might Quit*, WALL STREET JOURNAL (Mar. 13, 2015).

[70] *Id*.

[71] *Hadley*, 156 Eng. Rep. 145.

Second, the monitoring made possible by the IOT that feeds big data algorithms can help to reveal *ex post* hidden actions, which would facilitate contracting on *actual* effort, rather than expected effort or observable outcomes, again ameliorating moral hazard. For example, rather than adjusting premiums based on proxies for risk like age, education, and zip code, rates could be adjusted monthly or weekly based on actual driving results—speed, distance, locations, times.[72] Similarly, wearable sensors that transmit vitals could be used by health insurers to monitor activities that impact expected future payouts, such as exercise or alcohol consumption. Contracts, again, could be based directly on effort taken to avoid a medical incident rather than proxies for the probability that one will occur.

* * *

Adverse selection and moral hazard are real problems that force good types to subsidize bad types and reduce resources available to society as whole. Big data has the potential to ameliorate these problems by revealing more granular distinctions within distributions. That said, to ameliorate information asymmetries necessarily implicates privacy. In the next part, I develop a framework that helps distinguish productive from dissipative uses of big data.

III. WHEN IS BIG DATA HARMFUL?

As seen in Part II, big data is a potential salve to market failures arising from asymmetric information. But to determine that big data creates efficiencies is not to prove that restrictions on big data are harmful. It's trivial to claim that more information improves efficiency; privacy is valuable too, and in some contexts probably is so valuable that society is willing to forgo big data benefits to preserve privacy. What's more, some information collection can itself be dissipative. In this section, I draw on the economic theory of contracts and torts to develop a framework that identifies instances of dissipative privacy and concealment and also provides a mechanism for balancing big data against privacy when both information collection and concealment add value.

A. FRAMEWORK

Suppose that consumers—used collectively to refer to potential customers, borrowers, employees, or insureds—have private information (R)

---

[72] Some auto insurers, for example, already offer consumers the option to have their rates determined by car-based sensors that track a variety of metrics correlated with the probability of an accident.

about themselves. Consumers with higher *R*s are better types (e.g., better credit or insurance risks, or better employees). Firms (collectively referring to insurers, employers, lenders and producers) can use big data to reveal consumers' *R*s at a cost of *c*. Firms generate revenue from uncovering *R* through two channels. First, possessing information about consumers gives them an advantage in bargaining (e.g., by knowing reservation prices), which directly increases revenue by transferring *t* from consumers. Second, to incorporate the social value of revealing private information, I assume that firms also can use *R* to create surplus, *V*, by taking an action *x* (with a marginal cost of 1) that is customized to each value of *R* in the following manner: $RV(x) - x$. For example, this action may be more efficient matching of salaries with abilities, which would reduce wasteful expenditures on signaling or inefficiencies from pooling. Thus, firms' profits can be written as: $\pi = t - c + RV(x) - x$, where the first two terms are the net gain from using big data to transfer surplus, and the last two terms are the net gains to society from better matching. From the outset we can observe that if big data does not result in any direct transfer of surplus (*t*), firms would never attempt to ferret out information unless it created value.

If policy allows firms to use big data to estimate *R*, they will take a unique action, x_i^* for each unique value of *R*. It can be shown that profit-maximizing level of *x* is positively related to *R*, so that firms take higher level of effort for relatively "good" types.[73] For example, those with higher *Rs* will receive better credit or employment offers than bad types, or low-valued users who previously were priced out of a market will receive coupons that draw them into the market. Alternatively, if regulation prevents the use of big data to differentiate types, firms take one action, $\bar{x}$, for everyone, based on the average type, $\bar{R}$.[74] These two scenarios are shown in Figure 1. The horizontal axis measures *R*, and the vertical axis measures *x*. The function *x**(*R*) shows the profit-maximizing action taken for each type, and the line $\bar{x}(\bar{R})$ shows the single level of action taken when types remain unknown. Because $\bar{x}$ is profit-maximizing only for $\bar{R}$, *V*(*x*) is lower than it could be when firms must take $\bar{x}$ for $R \gtrless \bar{R}$. Accordingly, the gap between *x**(*R*) and $\bar{x}(\bar{R})$ represents the social loss from pooling versus separating equilibria.

---

[73] The first order condition for profit maximization requires that $R\pi_x(x^*(R)) - 1 = 0$. Differentiating this expression with respect to *R* yields: $\frac{\partial x*}{\partial R} = \frac{-\pi_x}{R\pi_{xx}} > 0$. This framework is similar to that used in Steven Shavell, *Acquisition and Disclosure of Information Prior to Sale*, 25 RAND J. ECON. 20 (1994).

[74] It can be shown that when the exact level of R is unknown, but the distribution of *R* is known, the optimal action is to take the action that maximizes profits at $\bar{R}$ for all types.

Figure 1:
Optimal Action by Type

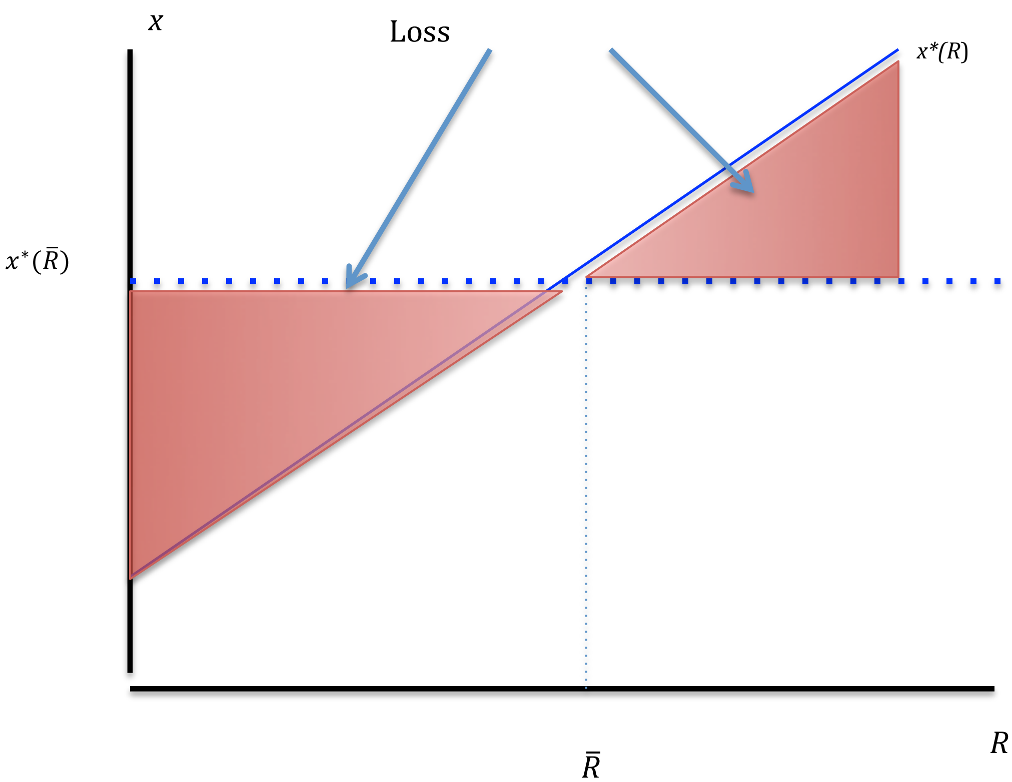

Consumers receive payoff: $U = x(R)+P - t$. Utility increases in x, which represents more favored actions due to higher Rs (e.g., lower interest rates, greater credit limits, lower insurance premiums). Loosely, x can be thought of as a consumer's share of the surplus generated through a firm's action. Consumers also receive P from the ability to conceal certain facets of themselves. This is the intrinsic value of privacy; people clearly value being free from unwanted observation, although this value varies across the population and contexts.[75] Relatedly, there is also a social value to privacy in the sense that forced revelation can reduce incentives to engage in productive activities—a sort of inverse moral hazard that underlies the theory of privileges that attach to conversations between doctors and patients, attorneys and clients, and husbands and wives. Even if it prompts better insurance or employment matching, for instance, revelation of HIV status may reduce incentives to become tested in the first place, although such knowledge clearly is valuable.[76] More generally, ubiquitous surveillance and predictions from the resulting data can lead to wasteful

[75] *See* notes 96- 108, *infra,* and accompanying text.
[76] *See, e.g.*, Benjamin E. Hermalin & Michael L. Katz, Privacy, *Property Rights and Efficiency: The Economics of Privacy as Secrecy*, 4 QUANT. MKT'G & ECON. 209, 212 (2006).

privacy protective behavior analogous to the wasteful expenditures on protecting property when property rights are ill-defined. For example, to avoid the consequences of being predicted to be at risk for diabetes, one may attempt to conceal their suspect grocery purchases, such as by purchasing sugary foods with cash.[77]

Consumers lose t, which is a transfer to firms with big data-driven knowledge of their reservation prices. Comparing a regime of privacy (in which firms choose $\bar{x}(\bar{R})$), to a regime of information revelation, in which a firm adopts an action tailored to each consumer, it is easy from to see in Table 1 how bad types lose in a separating equilibrium.

Table 1: Consumer Payoffs in Privacy and Revelation Regimes

		Legal Regime	
		Revelation	Privacy
Consumer Type	Good	$x_G{}^*(R_G) - t$	$\bar{x}(R_G) + P$
	Bad	$x_B{}^*(R_B)$ - t	$\bar{x}(R_B)$+ P

The impact of big data revelation is unambiguously negative for bad types $(R_B < \bar{R})$. They gain from revelation only if $(x_B^* - \bar{x})$ > P + t, which can never hold because $(x_B^* - \bar{x}) < 0$. Thus, regardless of intrinsic privacy concerns, there is clearly an incentive for bad types strategically to prevent algorithms from predicting their type. Good types$(R_G > \bar{R})$ may prefer revelation or privacy regimes depending on whether $(x_G^* - \bar{x}) \; > P + t$. Thus, whether

[77] This genre of privacy harm also meshes well with rights-based notions of privacy, which—apart from losses in utility from embarrassment that comes with disclosure of personal facts— focuses on notions of autonomy that are necessary to spur the type of creativity that serves society as a whole. *See, e.g.*, Joel Reidenberg, *Privacy Wrongs in Search of Remedies*, 54 HASTINGS L.J. 877 (2003); Daniel Solove, *Introduction: Privacy Self-Management and the Consent Dilemma*, 126 HARV. L. REV. 1880, 1892 (2013). For example, Julie Cohen argues that "lack of privacy means a reduced scope for self-making," and will shrink "the capacity for democratic self-government." Cohen, *What Privacy is For*, 126 HARV. L. REV. 1904, 1911 (2013). Similarly, Neil Richards has developed the notion of "intellectual privacy," which is autonomous space that is needed to develop ideas that have social value. Neil Richards, *Intellectual Privacy*, 87 TEX. L. REV. 387, 407 (2008). A rights-based inalienability rule for some personal information also could be justified in a utilitarian context if revelation sufficiently reduces incentives to create information in the first place, or increases incentives to invest in dissipative concealment.

good types prefer privacy to revelation depends on their intrinsic value of privacy relative to their share of increased in surplus from reductions in adverse selection, net of any pure transfers to firms.[78]

To ameliorate the impact of their type on their share of surplus, consumers can also take an action, y—with marginal cost of δ— that improves their value of R. For example, as discussed in Part II, insurance contracts based on true risk can spur consumers to drive more safely or adopt healthier living habits. In this manner, y captures the marginal reduction in moral hazard due to big data. Consumers have incentives to take actions to improve their types when the legal regime allows the use of big data to discover values of R. Consider the following game in which a consumer must choose between taking action, y, or no action.

Table 2: Moral Hazard with Privacy and Revelation

		Legal Regime	
		Revelation	Privacy
Consumer Action	y	$x^*(R+y) - \delta - t$	$\bar{x} - \delta + P$
	None	$x^*(R) - t$	$\bar{x} + P$

A consumer's dominant strategy is to take no action in a privacy regime, as taking y has no impact on x and will cost $-\delta$. On the other hand, as long as it is not too costly to engage in y ($\delta < (x^*(R+y) - x^*(R))$), the dominant strategy without privacy is to take action, which increases a consumer's payoff.[79] Importantly, taking y also improves welfare because it raises the average level of x by shifting the distribution of R to the right, which in turn increases $V(x)$.

In what follows, I use this framework to identify instances in which regulation of big data is likely to be welfare-enhancing.

[78] In a more general model, the tradeoff would also depend on the consumer's marginal rate of substitution between the economic benefits from revealing personal information (R) and the intrinsic privacy harms from such revelation (P): $\frac{\partial U/\partial x}{\partial U/\partial P}$.

[79] If δ is distributed throughout the population rather than constant, it may be the case that only some proportion of the population finds it cost effective to engage in y.

B. DISSIPATIVE CONCEALMENT AND REVELATION

Privacy is valuable intrinsically, but it's also valuable when concealment of relevant information leads to better terms of trade—"strategic concealment." A key difference between these types of privacy, however, is that value is created when concealment of information satisfies a demand, whereas strategic privacy concerns only securing a larger share of available surplus, not creating it. Accordingly, strategic privacy is purely dissipative; while it's privately rational to want to conceal information that will reduce your share of surplus from a bargain, such privacy is socially wasteful. Accordingly, to the extent that big data-driven separation thwarts strategic privacy, it should be counted as a benefit rather than a harm. At the same time, it has long been know that there can be socially excessive incentives to collect information; over forty years ago, Hirshliefer showed how investment in foreknowledge of events to gain a trading advantage is pure social waste unless the public revelation of this information spurs some surplus-creating action.[80] Otherwise, knowledge serves only to redistribute surplus, and expenditures to collect it are dissipative. Any sensible policy toward big data, therefore, should attempt to avoid promoting privacy or disclosure that serves *only* to move surplus from one party to another. That is, it is essential to distinguish between valuable and dissipative information collection and concealment.

A paradigm for this framework can be found in contract law. For example, sellers typically are required to disclose unfavorable information about their wares.[81] The rationale is that buyers will be able to make productive use of this information—to allow concealment would be to squander surplus for the seller's private gain.[82] On the other hand, buyers generally have no duty to disclose productive information that they have garnered, and for good reason; absent a property right to their information, buyers would have muted incentives to cultivate it in the first place, and again society would be worse off.[83] At the same time contract law tends to encourage the creation of productive information, it discourages investment in information that merely transfers surplus, such as insider trading or

---

[80] Jack Hirshleifer, *The Private and Social Value of Information and the Reward to Inventive Activity*, 61 AM. ECON. REV. 561 (1971)..

[81] MICHAEL J. TREBILCOCK, THE LIMITS OF FREEDOM OF CONTRACT 114 (1997); Steven Shavell, *Acquisition and Disclosure of Information Prior to Sale*, 25 RAND J. OF ECON. 20 (1994).

[82] *See* STEVEN SHAVELL, FOUNDATIONS OF ECONOMIC ANALYSIS OF LAW (2006).

[83] For example, if the buyer has information about mineral deposits on land, he has no duty to disclose. ALEX M. JOHNSON JR., UNDERSTANDING MODERN REAL ESTATE TRANSACTIONS (3d. ed. 2012).

foreknowledge of a conditions that impact the value of a commodity.[84] The distinction between duress and necessity also has an economic rational that rests on the distinction between creative and dissipative actions. Allowing recovery for bargains made under duress would encourage resources devoted to trying to wrest surplus from others, and concomitant expenditures to defend these attempts.[85] Allowing bargains made out of necessity to stand encourages the supply of value-enhancing aid, and the limitation on consideration mutes incentives to over-invest in rescue.[86] Finally, the limitation on consequential damages creates incentives for buyers to reveal private information about their sensitivity to breach.[87] Concealment in these circumstances in wholly dissipative, as it forces normal types to subsidize sensitive types. These doctrines are all designed to reduce incentives to spend resources to merely transfer wealth, and can provide a blueprint for distinguishing productive from dissipative privacy.

1. *Dissipative Concealment*

For privacy to be dissipative, three conditions must be met. First, concealment of R must retard value-creating actions—actions that reduce adverse selection (x) or moral hazard (y). That is, knowledge of R leads to higher values of $V(x)$. Second, the value created from these actions must be greater than their marginal cost ($\frac{\partial x*}{\partial R}V' > c$). Finally, there must be no intrinsic privacy gain from concealment ($P = 0$). If these conditions are met, privacy is dissipative because the only gains from concealment come in the form of increased share of surplus to bad types ($(x_B^* - \bar{x})+ t$), and at the expense of total surplus. This is easily shown. For example, consider a world with two workers—one good and one bad. If there are no intrinsic privacy gains from concealment, privacy is welfare enhancing only if:

$$(\bar{x} - x_B^*) > (x_G^* - \bar{x}) + [(R_G(V(x_G^*) - V(\bar{x})) - x_G^* + \bar{x}) + (R_B(V(x_B^*) - V(\bar{x})) - x_B^* + \bar{x})].$$

The left-hand side of this expression is the gain to bad types from concealment. The first part of the right-hand side is the gain to good types from revelation, and the second part is the gain to society from increased surplus due to better matching. The second part of this expression is always

[84] *See Laidlaw v. Oregon*, 15 U.S. 178, 194; Shavell, *supra* note 71, at __.
[85] *See* ROBERT COOTER & THOMAS ULEN, INTRODUCTION TO LAW AND ECONOMICS (6th ed. 2011).
[86] *Id.*
[87] *See Hadley*, 156 Eng. Rep. 145; Ian Ayers & Robert Gertner, *Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules*, 99 YALE L.J. 87 (1989).

positive because $RV(x^*)-x^* \geq RV(\bar{x})-\bar{x}$ for all types.[88] So the condition for purely strategic privacy to be socially beneficial can be reduced to the following necessary (but not sufficient) condition:

$$\bar{x} > \frac{x_G^*+x_B^*}{2},$$

which can never hold because $\bar{x} = \frac{x_G^*+x_B^*}{2}$.

To make this result more concrete, consider the potential lazy employee whose productivity score from an accurate big data algorithm is too low to garner an interview. The prediction that he's unsuitable for employment is disappointing to the applicant because he is no longer able to cloak his true type. That is, the only harm from revelation is the loss of $(x_B^* - \bar{x})$. And because the ability to sort good from bad workers raises productivity ($\frac{\partial x*}{\partial R} > 0$), losses to bad types from separation are less than the gains to society as a whole, which include increased firm profits and increased utility to those with high values of R whose market opportunities previously were limited due to adverse selection.[89] Put differently, when privacy serves purely strategic purposes, losses to bad types due to big data-driven sorting should never be counted as privacy harm because they are merely artifacts of a net social benefit due to a reduction in adverse selection; without these losses, the net gains to society cannot materialize.[90] Here, information is put to its most valuable use when it's revealed.

2. *Dissipative Revelation*

If big data predictions do not prompt surplus-enhancing actions (*i.e.*, $\frac{\partial x*}{\partial R} = 0$), privacy can never be dissipative. Information collection that has no impact on surplus, therefore, is purely dissipative: firms are spending c to transfer t from consumers to themselves.[91] This is true whether or not consumers gain intrinsic value, P, from concealment; when information collection costs more than the value it creates, privacy is always the most efficient policy as it preserves value. It prevents expenditures that merely transfer surplus.

---

[88] This is because by assumption x^* maximizes $R(V(x)-x)$.

[89] These gains may also include reductions in moral hazard from choosing to engage in y.

[90] An analogy can be found in the *per se* condemnation of naked agreements among firms to fix prices, allocate markets, or otherwise to compete less vigorously. Although such agreements are privately beneficial to their participants, they unambiguously reduce social welfare. Accordingly, the antitrust laws do not countenance any defenses to *per se* conduct. *See* Thomas G. Krattenmaker, *Per Se Violation in Antitrust Law: Confusing Offenses with Defenses*, 77 GEO. L.J. 165 (1988).

[91] Recall that if data collection provides no increase in value, it would be rational to collect information only if $t > 0$.

Consider the following hypothetical: a job-matching algorithm that predicts only sexual orientation. A firm will not be able to use this information to create surplus because it has no bearing on productivity ($\frac{\partial x*}{\partial R} = 0$). Thus, infringing on privacy in this manner has no social value, and categorizing this type of prediction as a privacy harm serves to preserve surplus in two ways: discouraging firms from expending c in discovering this type of information, and eliminating any direct privacy harms (-P) from its revelation. A similar case can be made for the use of big data to create "sucker lists" of vulnerable consumers. This investment creates no social value and, like expenditures on rent-seeking, serves only to transfer surplus (t) from consumers to firms at a cost to society of c. Here again, information is most valuable to society when it remains concealed.

C. CONCEALMENT AND REVELATION BOTH VALUABLE

Things become more complicated when the discussion turns to sensitive data with both intrinsic and strategic value. In these cases, neither revelation nor privacy are dissipative: firms can increase surplus with big data predictions, but consumers also gain P from concealment. In this section, I examine factors that suggest presumptions in favor of revelation or concealment.

1. *Factors Influencing Gains from Separation*

In Figure 2, there are two possible $x(R)$ functions: x^*_L and x^*_H. Recall that in a privacy regime firms take a uniform action, $\bar{x}(\bar{R})$, for all types of R, so that the welfare loss from adverse selection can be represented by the gap between $\bar{x}(\bar{R})$ and the optimal action with respect to type R_i ($x_i^*(R_i)$). It's easy to see that welfare losses from privacy are larger the steeper is x^*, reflecting the fact that optimality requires a large variance in action over types. Further, these losses will be exacerbated by moral hazard because incentives to take action y is a positive function of the gap between $\bar{x}(\bar{R})$ and $x_i^*(R_i)$ in a revelation regime.

The total social costs associated with the gap between $\bar{x}(\bar{R})$ and $x_i^*(R_i)$ are also determined by the density of the population at each R.[92] In Figure 2, there are also two possible distributions of R, f_1 and f_2, with the latter being more dispersed, representing a more heterogeneous population. If most of the population is centered around $\bar{R}$ (distribution f_1), the welfare losses from privacy associated with both x^*_L and x^*_H will be similarly small,

[92] Formally, social losses are: $\int_{\underline{R}}^{\overline{R}} [V(x^*(R_i)) - V(\bar{x}(\bar{R}))] f(R) dR$.

as for the vast majority of the affected population the gap between both x^*_L and x^*_H and $\bar{x}$ is relatively small. Thus, even if there are large gains from separation, if most of the population is homogenous over the trait of interest, the inefficiencies from pooling will be small. Alternatively, if the population varies over the trait in interest, R will be distributed more like f_2. Because a much smaller proportion of the population is centrally located, that there will be non-trivial social losses for pooling even if the gains from separation are more akin to x^*_L than x^*_H.

Figure 2:
Losses from Mismatch of Optimal Action

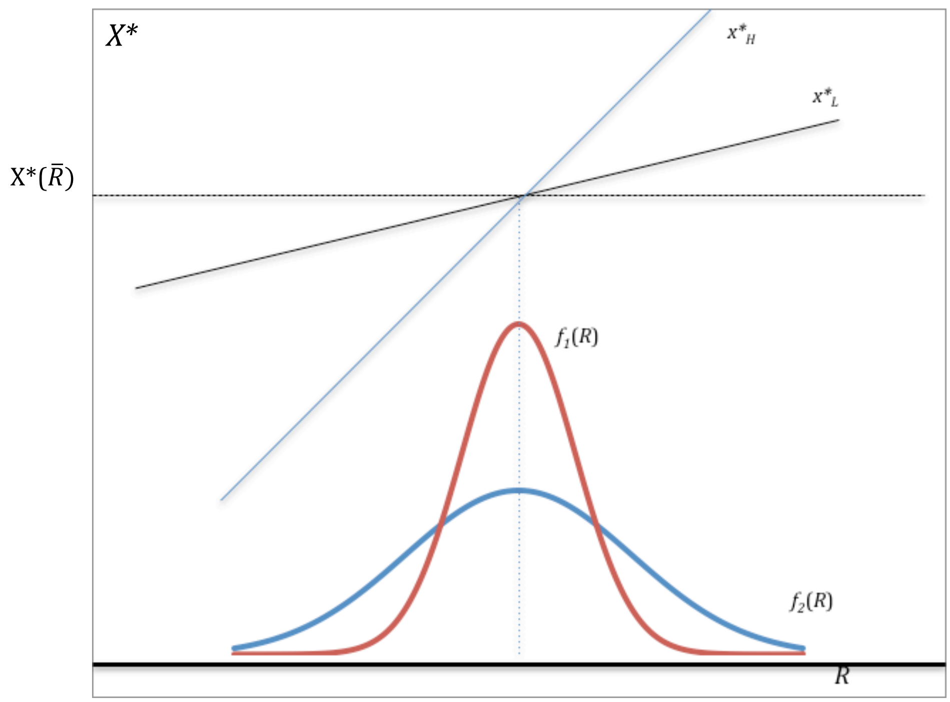

Consider the following example. Suppose that R measures ability successfully to complete law school and x is the discount on tuition. There are likely to be large gains from matching types to tuition, so that $x^*(R)$ is relatively steep; those who are unlikely to succeed (low Rs) will be discouraged from attending and wasting their time and money, whereas those with high aptitudes for a legal career will be encouraged to acquire legal training. If law schools were barred from collecting data to discover abilities (*e.g.*, through requiring the LSAT or undergraduate grades, or big data algorithm that relied on non-traditional data), they would offer an average tuition based on the average quality of the pool they expect to attract. This rate, however, would attract some who will not complete the program and discourage some who would be quite successful. Although the social costs associated with mismatches may be large, if the pool of applicants is relatively homogenous, the incentives to attend law school will be approximately optimal for most of the population—only those few at the extremes of the distribution have severely distorted incentives. On the other hand, if applicants are quite diverse over their ability to complete law school, the losses from the most severe mismatches receive more weight.

2. *Identifying Optimal Privacy Regulation*

Having identified circumstances in which gains from revealing private information through big data are likely to be large or small, we can marry this framework with the standard economic model of accidents to gain some insight into when retarding big data may be appropriate. In the standard model, the optimal level of care is found by minimizing the sum of accident and avoidance costs:[93]

$$P(z)^{*}H + \theta z.$$

$P(\bullet)$ is the probability of a privacy accident—some sort of big data prediction that causes intrinsic privacy harm, H. A firm can take an action, z, to reduce the likelihood of harm, which in the context of big data includes retarding the collection of certain data, the use of data to make certain predictions, or certain uses of predictions. Retarding big data practices reduces the likelihood of a privacy harm, but comes at the social cost identified above, which will vary depending on the distribution of types and the gains from separation as discussed in the previous section. Further, social costs will also vary by regulation type: restrictions on collection will entail larger costs than restrictions on use, as the former type of regulation not only entails substantially larger direct costs (e.g., notice and consent mechanisms), but also eliminates all possible future uses of the data, some of which may be beneficial. These cost are represented by θ.

In Figure 3, there are two curves, z_L and z_H, which map optimal levels of care for various levels of harm—clearly, the higher the harm, the more the care.[94] The differences between theses curves are the benefits from big data driven separation.[95] For z_L, the benefits from separation are small, so retarding collection or use of data comes at little cost (θ_L) . The marginal

---

[93] *See* Shavell, *supra* note 71, at 177-79.

[94] The curves are concave because of the diminishing marginal effectiveness of additional precautions: $\frac{\partial z*}{\partial H} = -\frac{P'}{HP''}$.

[95] It can be shown that $\frac{\partial z*}{\partial \theta} = -\frac{1}{HP''} < 0$.

costs and benefits from care are equal at H^*_L ($z_L(H) = 0$). For harms less than H^*_L, retarding information flows leads to net harms from strategic concealment, which implies that the optimal level of regulation is zero. For harms greater than H^*_L, positive levels of care (*i.e.*, some form of concealment) are socially justified because intrinsic harms are greater than the benefits from revelation. Alternatively, for z_H, the benefits from separation are relatively large (θ_H), suggesting that the optimal level of privacy regulation is zero until a higher threshold level of harm, H^*_L is met.

Figure 3:
Optimal big data Restrictions with Heterogeneous Harms and Benefits

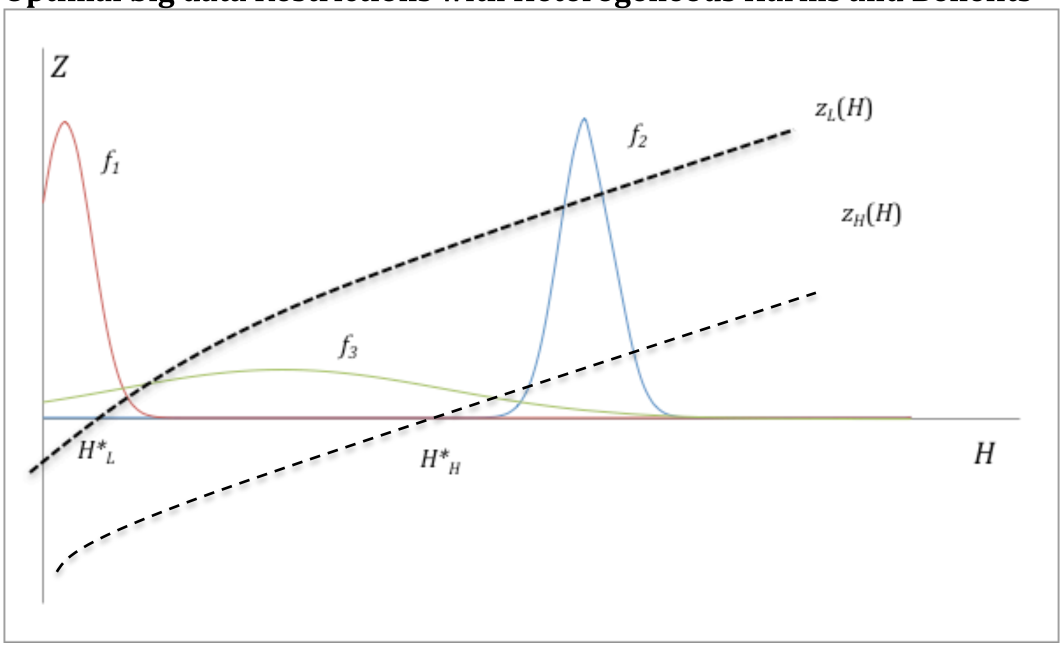

Because intrinsic privacy harms are felt differently across a population, H is distributed as $f(H)$. Consider three distributions of intrinsic privacy harm shown in Figure 3. The first, f_1, shows the distribution of harm from a practice that most consider innocuous: it is truncated at zero, and it is dispersed, with the tail representing the presence of a small number of privacy sensitive people. The distribution f_2, on the other hand, represents the harm associated with the revelation of information that most people agree is highly sensitive; the average harm is large and the variance is small. Finally, f_3 is a distribution of sensitivity to information that reflects a heterogeneous population—intrinsic harm ranges from relatively low to relatively high. Unlike the harm in f_1, even those with the lowest sensitivities suffer some intrinsic privacy harms. At the same time, only the tail of the distribution suffers the level harm associated with f_2. Further, because the distribution of harm is so broad, the "average" level of harm is of little significance—that is, unlike for f_1 and f_2, one cannot approximate the level of harm for most of the population with the mean.

All of this underscores how crucial it is to have information about intrinsic privacy harms to arrive at sensible policies in this area; to weigh intrinsic privacy harms against big data gains, one must be able to quantify them. This exercise is far easier said than done. It's relatively straightforward to quantify losses from identity theft or credit card fraud, which are objectively measured in dollars that (leaving aside diminishing marginal utility of wealth) do not vary in value across the population. These losses, however, are not the types of privacy concerns typically associated with big data. To determine when limiting big data may make sense, we must answer questions like how much do people value preventing algorithms from predicting health status, income, driving ability, taste in clothes, or sexual orientation? What about sharing anonymous geo location histories, web surfing habits, or social media posts with faceless servers? Concealment of these facts surely has intrinsic value to some, but these values are highly subjective, which renders them unverifiable against objective measurement. Making this exercise even harder, privacy values are likely to vary across populations and contexts.

The available empirical evidence provides little guidance. For example, survey data show that consumers care about privacy, yet revealed preferences suggest stated concerns may be exaggerated. In a recent Pew Poll, 65 percent of respondents say that "controlling what information is collected about you" is "very important."[96] At the same time, consumers increasingly participate in online activities that reveal person data to known and unknown third parties; the percentage on online adults engaging in social media rose from 8 percent in 2005 to 72 percent in 2013.[97] And the health tracking market has exploded in recent years.[98] At the same time, very few people bother to opt-out of online tracking or adopt privacy-protecting technology, like the TOR browser or searching via Duck, Duck, Go![99] Indeed, Acquisti, Taylor, and Wagman conclude in a recent survey of the literature that the adoption of privacy enhancing technologies has lagged substantially behind the use of information sharing technologies.[100] Thus,

---

[96]Mary Madden & Lee Rainie, *Americans' Attitudes About Privacy, Security, and Surveillance*, PEW RESEARCH CENTER, at 5 (May 15, 2015), http://www.pewinternet.org/2015/05/20/americans-attitudes-about-privacy-security-and-surveillance/.

[97] Joanna Brenner & Aaron Smith, *72% of Online Adults are Social Networking Site Users*, PEW RESEARCH CENTER, at 2-3 (Aug. 5, 2013), http://www.pewinternet.org/2013/08/05/72-of-online-adults-are-social-networking-site-users/.

[98] Susannah Fox, *The Self-Tracking Data Explosion*, PEW RESEARCH CENTER (June 4, 2013), http://www.pewinternet.org/2013/06/04/the-self-tracking-data-explosion/.

[99] See Maurice E. Stucke & Allen P. Grunes, *No Mistake About It: The Important Role of Antitrust in the Era of Big Data*, ANTITRUST SOURCE at 8-9 (April 2015).

[100] *See* Alessandro Acquisti et al., *The Economics of Privacy*, J. ECON. LIT. at 37-38 (forthcoming, 2016), available at (SSRN).

although consumers tell survey-takers that they are concerned about privacy, consumers' revealed preference suggests that privacy concerns are not sufficient to slow the adoption of services that rely on the collection and use of their data data.

Further, recent work by Benjamin Wittes and Jodie Liu suggests that people are more privacy concerned with proximate observation by individuals than distant observation by computers.[101] They note that commentators tend to ignore the privacy benefits that come from the ability to find and consume information or goods in private. For example, they find evidence from Google Autocomplete that people often search for information on HIV and sexual identification, suggesting that the ability to search anonymously online for information about these topics provides an important privacy benefit and probably spurs increased information generation. Research in a similar vein finds that self-checkout in libraries has increased the number of LGBT books checked out by students, again suggesting that privacy concerns are reduced when human interaction is removed from the situation.[102]

There has been some work at trying to measure intrinsic privacy valuation, but the thin extant literature provides little that is generalizable and is. For example, a series of papers by Alessandro Acquisti and various co-authors uses experimental methods to test whether subjects suffer from various cognitive biases when making decisions about privacy. The authors find that consumers appear to value privacy more when they are asked to sell it than when they must purchase it.[103] Consumers' willingness to divulge private information also appears to depend on context and cues.[104] Further, the perceptions about the ability to control one's information impact willingness to share personal information.[105] Other researchers have found that consumers would be willing to accept small discounts and purchase recommendations in exchange for personal data,[106] and that they exhibit low

---

[101] Benjamin Wittes & Jodie Liu, *The Privacy Paradox: The Privacy Benefits of Privacy Threats*, CENTER FOR TECHNOLOGY INNOVATION AT BROOKINGS (May 2015), http://www.brookings.edu/~/media/research/files/papers/2015/05/21-privacy-paradox-wittes-liu/wittes-and-liu_privacy-paradox_v10.pdf.

[102] *See also* Stephanie Mathson & Jeffry Hancks, *Privacy Please? A Comparison Between Self-Checkout and Book Checkout Desk for LGBT and Other Books*, 4 J. ACCESS SERVS. 27, 28 (2007).

[103] *See* Alessandro Acquisti, Leslie K. John, and George Lowenstein, *What is Privacy Worth?,* 42 J. LEG. STUD. 249 (2013).

[104] *See* Alessandro Acquisti, Leslie K. John, and George Lowenstein, *Strangers on a Plane*, 37 J. CONSUMER RES. 858 (2011).

[105] *See* Laura Brandimarte, Alessandro Acquisti, & George Lowenstein, *Misplaced Confidences: Privacy and the Control Paradox*, 4 SOC. PSYCHOL. AND PERSONALITY SCI. 340 (2012).

[106] *See* Dan Cvrecek, Marek Kumpost, Vashek Matyas & George Danezis, *A Study on the Value of Location Privacy*, Proceedings of the 5th ACM Workshop on Privacy in the Electronic Society (2006). *See also* Acquisti et al, *supra* note 89, at 39 for a review of the empirical literature.

willingness to pay to for protection from telemarketers.[107] For example, one study finds that consumers are willing to pay an additional $1-$4 for a hypothetical smartphone app that conceals location, contacts, text content, or browser history from third-party collectors.[108]

The point here is not that consumer valuation of privacy shouldn't count because it cannot be quantified. To the contrary, subjective harms are real and optimal deterrence should take account of them. Nonetheless, given the current state of knowledge their measurement is little more than guesswork. Given the costs associated with over deterrence of beneficial practices, policy makers should proceed with caution. When policy makers measure harms inaccurately, they may retard beneficial information flows. Regulatory responses to worst-case hypotheticals or demands from the most privacy sensitive can do more harm than good by forcing consumers to suffer the ill-effects of adverse selection and moral hazard facilitated by strategic privacy. Further, inaccuracy creates uncertainty for businesses trying to comply with the law. If businesses are unsure about where the line between legal and illegal behavior is drawn—which is a function of the estimated distribution of harm—they rationally will take too much care to avoid violating the law.[109] In the context of big data, "too much care"—can be mean self-limiting beneficial uses of data. When objective measures of harm are absent, moreover, policy makers (e.g., Federal Trade Commission Commissioners and staff) are able to import their own subjective judgments of harm to guide policy. For example, the FTC's recent report on data brokers identified as potential harm from an array of hypotheticals, including "being limited to ads for subprime credit," and the facilitation of the "sending of advertisements about health, ethnicity, or financial products, which some consumers may find troubling."[110] Based on these "findings", the FTC recommended imposing FCRA-like requirements on data brokers involved in non-FCRA activities.[111] In this instance, there is no indication of harm, let along attempt to quantify it. Instead, policy recommendations seemingly are based on various Commissioners' views of what they find potentially

---

[107] *See* Hal R.Varian, Glenn Woroch & Fredrik Wallenburg, *Who Signed Up for the Do Not Call List? (2004), http://eml.berkeley.edu/~woroch/do-not-call.pdf*; Ivan P. L. Png, *On the Value of Privacy from Telemarketing: Evidence from the "Do Not Call" Registry* (2007), at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1000533.

[108] Scott Savage & Donald M. Waldman, *The Value of Online Privacy* (2013), *at* http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2341311.

[109] This is a well-known result in the economics of accidents, and is due to the discontinuity in total costs—accident and avoidance—at the negligence standard. *See* Kostad *et al*., *Ex Post Liability vs. Ex Ante Safety Regulation: Substitutes or Complements*, 80 AM. ECON. REV. 888, 894-95 (1990); Shavell, *supra* note 71, at 224-29.

[110] Federal Trade Commission, *supra* note 42, at 48.

[111] *Id.* at 55-56. and the Chairwoman and Commissioner Brill recommended requiring that data brokers assure that their data sources obtained data "with notice and choice" and "express affirmative consent for sensitive data." *Id.* at 52 n.91.

offensive to privacy norms. When policy is based on subjective measures of harm, it expands agency discretion, which in turns enhances incentives to engage in rent seeking; spending on redistribution rather than production.[112]

Striking a precise balance between intrinsic demands for privacy and efficiencies from separation may be beyond the current state of knowledge. Nonetheless, the framework developed in this paper can help identify factors that should influence the proper default regulatory posture; that is, who bears the burden of proof—those proposing restrictions, or their opponents? Examining these questions in this framework has the advantage of relying on information that is more readily available than that on intrinsic privacy values to help inform proper regulatory postures. For example, the distribution of creditworthiness is generally known,[113] and as discussed in Part II, there is a large literature measuring the presence of adverse selection and moral hazard. These two pieces of known information can provide a threshold of harm that would be necessary to justify big data restrictions, and in turn can inform regulatory defaults.

It is likely to make sense to have a presumption in favor of big data, when the following conditions are present:

- Big data is aimed at separation on a trait that is relatively widely dispersed;
- Gains from reducing adverse selection and moral hazard are likely to be large;
- Predictions involve non-sensitive traits; or
- Predictions tell little about reservation prices.

This scenario can be represented by f_1 and z^*_H in Figure 3. When the gains from big data are large, the demand for regulation is non-existent, as there is no mass to the right of H^*_H. Even when the benefits from big data are relatively small (z^*_L), only a tiny fraction of the population—the mass in f_1 to the right of H^*_L—suffers harm at a level that justifies any care. In an ideal world, big data would be restricted just a little for those consumers, but in the real world in which choices are lumpy rather than continuous and usually cannot be tailored to individuals. One must decide on a default presumption, and in these cases the best policy posture is one of "permissionless innovation," in which the burden is on those advocating restrictions to show harm.[114]

---

[112] *See* James C. Cooper, *The Perils of Excessive Discretion: The Elusive Meaning of Unfairness Under Section 5 of the FTC Act*, 3 J. ANTITRUST ENFORCEMENT 87 (2014).
[113] *See*, *e.g.*, http://www.fico.com/en/blogs/?s=distribution+.
[114] *See* Adam Thierer, *Permissionless Innovation: The Continuing Case for Comprehensive Technological Freedom*, MERCATUS CENTER, at 9 (2015).

So, what types of practices likely fall into this bin? Analytics used for online and offline marketing is one candidate. The observation and analysis of online and offline shopping habits by an anonymous algorithm to make predictions about the types of goods and services one likes involves the collection of, and predictions relating to, relatively non-sensitive information. Further, even if the gains from identifying distinct tastes and preferences are not likely to be trivial.[115] This scenario also would cover instances where the vast majority of gains from privacy are strategic. For example, using driving data or credit scores for auto insurance involves a relatively non-sensitive prediction (driving ability), and the gains in terms of separation—both reduction in adverse selection and moral hazard—stand to be large. Likewise, the use of social media posting and other unconventional data streams for alternative credit scoring also are likely to provide large separation gains, and creditworthiness is not typically considered sensitive information. Although some good types may be so privacy sensitive that they prefer to forego the gains from separation, most of the gains from forced pooling here are likely to be strategic in nature. What's more, the predictions here are geared toward pricing different risk profiles rather than discerning reservation prices for similar risk profiles to gain a larger share of the surplus.

The following characteristics militate toward a more aggressive regulatory stance:

- Homogeneous populations over the trait in questions;
- Little problem with adverse selection or moral hazard;
- Predictions involving highly sensitive traits; or
- Predictions that are geared toward discerning reservation prices.

For example, if privacy harm can be represented by f_2, even if big data offers large efficiencies (z^*_H) the entirety of the population suffers harm sufficient to justify some form of restrictions (to the right of H^*_H). Distributions like f_2 may be privacy harms involving sensitive health status or children. The social benefits from regulation are even more clear when the benefits from big data are small, for example as would be the case for an algorithm that predicted the presence of a rare genetic disorder for which there was no treatment. Although this prediction would result in more efficient *ex post* matching for insurance purposes, it would not reduce moral hazard and any gains from reducing adverse selection likely would be trivial (i.e., $\frac{\partial x*}{\partial R} \cong 0$). Further, allowing these predictions to be used could discourage discovery of

[115] *See* Avi Goldfarb & Catherine E. Tucker, *Privacy Regulation and Online Advertising*, 57 MGM'T SCI. 57 (2011).

this information in the first place, which could be useful in the hands of the sufferer. Similarly, using big data to create "sucker lists" is an example of expending resources merely to determine a class of people from whom surplus can more easily be extracted. This use creates no social value, and serves only to transfer wealth from the frail to unscrupulous companies.[116] A default in favor of big data restrictions in these cases would make sense.

Of course the hard cases occur when the distribution of intrinsic harm looks like f_3 —that is, when there is little agreement or little accurate information on how consumers suffer intrinsic privacy harms. But again, we can turn to the more easily known gains from separation to get an idea of how beneficial regulation may be. When there are small gains from separation (z^*_L), a large part of the distribution is better off with some form of regulation (those to the right of H_L*), and when the gains are relatively large, a minority of population (those to the right of H_H*) will benefit from some form of restriction. Thus, a regulatory default makes sense only when the gains from separation are likely to be small, and even then any regulation should be less stringent than when there is agreement that privacy harms are large (*i.e*, f_2), such as requiring an opt-out for the specific use.

The Target incident may fit into this category. Although as noted earlier it has become the chief example of big data gone bad, it's not so clear that limiting Target's ability to make predictions about the pregnancy status of its customers would stand up to scrutiny under the framework. Target used data from its baby shower registry—which provided it with a list of women with known due data—to analyze shopping habits, with the goal of being able to send offers to women in their second trimester.[117] The benefits could be substantial. For example, if the goods advertised were unit elastic, a five percent reduction in price (from a coupon) would increase consumer surplus by five percent plus an amount proportional to the pre-coupon sales.[118] Further, to the extent that Target's mailers included discounts on prenatal vitamins or other products that would improve prenatal health, the benefits are even larger. [119] Thus, there are clear benefits to separation here that would be lost if Target were kept from making predictions about

---

[116] *See* David C. Vladeck, *Digital Marketing, Consumer Protection, and the First Amendment: A Brief Reply to Professor Calo*, 82 GEO. L.J. *ARGUENDO*, 156, 162 (2014).

[117] Charles Duhigg, *How Companies Learn Your Secrets*, NEW YORK TIMES MAGAZINE (Feb. 16, 2012).

[118] These gains are magnified if the targeted goods were more price elastic. Total surplus on a linear demand curve increases in the following manner in response to a price reduction: $\Delta Surplus = \%\Delta P * \frac{Q\varepsilon}{2}$, where ε is the absolute value of the price elasticity of demand.

[119] Even without price reductions, mere advertisements are likely to increase demand for pre-natal vitamins and hence total surplus. *See*., *e.g*., Dhaval Dave & Henry Saffer, *Impact of Direct-to-Consumer Advertising on Pharmaceutical Price and Demand*, 79 SOUTHERN ECON. J. 97 (2012) (finding that direct-to-consumer broadcast advertising of prescription drugs increased demand by about 12 percent).

pregnancy, or even collecting the data in the first place.[120] The distribution of intrinsic privacy harm from having one's pregnancy predicted by an algorithm, moreover, is likely to look like f_3. Although the predictions have to do with marketing, because it concerns pregnancy, more people are likely to some privacy harm than the behavioral targeting that is represented by f_1. At the same time, second-trimester pregnancy status does is not rise the level of medical conditions that come with stigma or embarrassment depicted in f_2. Although some women may want to conceal their pregnancy—e.g., from unapproving parents or from current or perspective or current employers[121]— this is unlikely the case for most women.[122] What's more, even if you want to conceal your pregnancy from the world, a Target algorithm correctly predicting your pregnancy is not the same as having it revealed to the world. The pregnant teen in the Target story notwithstanding, the odds that receiving discounts from a store in the mail will tip off those from whom you wish to conceal your condition are slim.

The framework also can shed some light on the debate over when restrictions on data collection should be used in conjunction with a focus on harmful uses. The recent report on big data from the President's Council of Advisors on Science and Technology (PCAST Report), for example recommends that policy focus should be on "uses of big data" rather than "collection and analysis."[123] Others, however, have argued that collection restrictions should not be abandoned so readily. As FTC Chairwoman Ramirez has said, "[i]nformation that is not collected in the first place can't be misused."[124] The model suggests that collection restrictions will improve welfare in only very narrow circumstances. As noted above, restrictions on

---

[120] *See* Vladeck and Bedoya, *supra* note__, at 6 ("strong ex ante use limitations could have stopped Target from identifying pregnant women through their purchases.").

[121] Alissa Quart, *Why Women Hide Their Pregnancies*, New York Times (Oct. 7, 2012), at http://www.nytimes.com/2012/10/07/opinion/sunday/why-women-hide-their-pregnancies.html?_r=0.

[122] What's more, by the second-trimester, most women have outward signs of pregnancy, making it difficult to conceal even if they wanted to. *See Pregnancy Stages: Your Baby, Your Body*, Webmd, at http://www.webmd.com/baby/features/pregnancy-stages-baby-body.

[123] PRESIDENT'S COUNCIL OF ADVISORS ON SCIENCE AND TECHNOLOGY, REPORT TO THE PRESIDENT: BIG DATA AND PRIVACY: A TECHNOLOGICAL PERSPECTIVE, 49 (May 2014). The PCAST Report explainthat it may be nearly impossible to determine *ex ante* which type of data is personal, when any data could be transformed into "personal" data when mixed with other data sets. Further, most data is dual-use in nature—there are uses that may be harmful to privacy, but others that are perfectly benign or beneficial. See id. at 50 ("the information in big data that may raise privacy concerns is increasingly inseparable from a vast volume of the data of ordinary commerce, or government function, or collection in the public square.").

[124] Chairwoman Edith Ramirez, Federal Trade Commission, *The Privacy Challenges of big data: A View from the Lifeguard's Chair*, Keynote Address, Technology Policy Institute, Aspen Forum (Aug. 19, 2013). *See also* Comments of Alvaro Bedoya & David Vladeck, Center for Privacy & Technology at Georgetown Law Center, on big data and Consumer Privacy in the Internet Economy, (Aug. 5, 2014).

collection increase costs of care (θ) directly through costs associated with notice and consent requirements.[125] Further, because collection restrictions prevent *all* future uses, there are indirect costs associated with lost beneficial uses. In this manner, collection restrictions have the effect of shifting $z(H)$ down, which raises the threshold level of harm needed to justify regulation.

Figure 4:
Optimal Care with Use and Collection restrictions

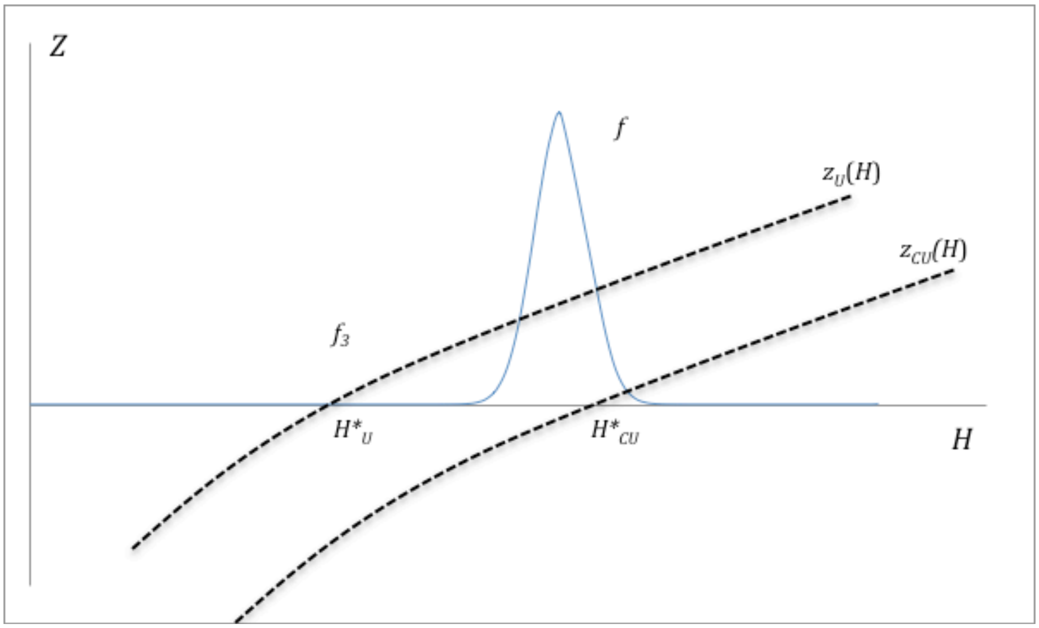

Figure 4 shows $z_U(H)$, which maps the optimal level of care for levels harm with use restrictions only, and $z_{CU}(H)$, maps optimal care for levels of harm with use and collection restrictions (e.g., notice and consent requirements or outright bans). Given the distribution of harm, all of the population suffers intrinsic harm that justifies use restriction (the full density is to the right of H_U*), but when regulation involves restrictions on collection, only the upper tail of the distribution suffers sufficient harm to justify regulation. This suggests that unless (1) there is broad agreement that harms are substantial for a large portion of the population, and (2) benefits are likely to be small for any possible use of this data, regulation should focus on use rather than collection.

IV. BIG DATA AND THE POOR

[125] The PCAST Report notes the potentially crippling expenses associated with enforceable collection regulation. PCAST Rep. at 50 ("The related issue is that policies limiting collection and retention are increasingly unlikely to be enforceable by other than severe and economically damaging measures.").

A common theme in the privacy literature is that big data disproportionately will harm the poor. As discussed below, however, most of these worries tend to fade when confronted with economic theory and empirical evidence. Indeed, the poor stand to gain more than the rich in many circumstances. Below, I use some of the insights developed in Part III to explore implications for the poor in three areas that have garnered a significant amount of attention in the big data debate: credit markets, price discrimination, and labor markets.

A. *Credit Markets*

Some have expressed concerns that so-called "data determinism," will marginalize the poor by limiting their options.[126] For example, in its recent report on data brokers, the FTC expressed concern that the poor will be marketed only subprime credit offers.[127] There are theoretical reasons to approach these hypotheticals with skepticism, and more importantly, empirical evidence to suggest that these concerns are likely to be misplaced.

From a theoretical perspective, there are reasons to expect the gains from big data separation to be larger for the poor than the rich. As sample size increases, a statistical test's power to detect small differences rises. This means that big data is likely to have the largest impact in intra-group separation. Take for example the distribution of credit worthiness in Figure 5. It doesn't take big data to figure out that consumers located at A and A' (prime market) are more likely to repay their loans than B and B'(subprime market)—such a large difference is easily estimated with confidence even with low power tests. Further, there is already likely to be separation between those at the top of the distribution because they interact often with the credit system; potential lenders will have an easier time estimating the difference between A and A' simply because there is already sufficient data on actually probability of repayment from repeated transactions. Those at the bottom of the distribution, however, have few interactions with the credit system, so they are largely unknown.[128] Accordingly, lenders have no choice but to judge them based on observables, such as address, education, or income —which don't vary between B and B'—even though underlying likelihood of default may vary substantially. All of this means that there is insufficient power to identify the true differences between B and B'; absent big data, B' is grouped with B despite the fact that she is less likely to default.

---

[126] *See* Ramirez, *supra* note 3, at 4.
[127] *See* Federal Trade Commission, *supra* note 42, at 48; Ramirez, *supra* note 3, at 4.,
[128] *See* note 142, *infra*, and accompanying text.

With big data, however, analytics can make finer predictions that allow B' to separate from B and receive better terms.

Figure 5: Distribution of Creditworthiness

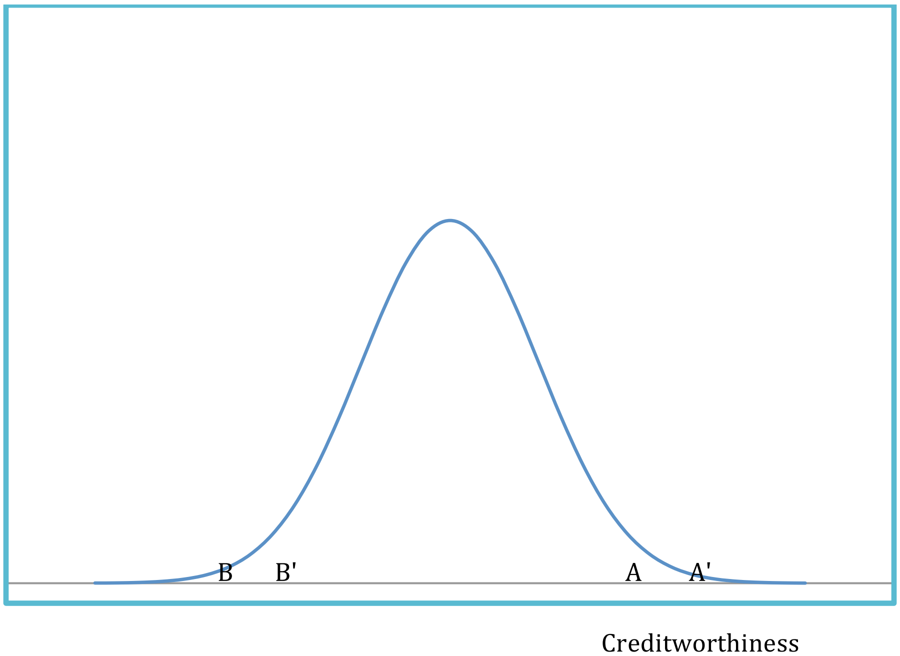

The marginal value of separation for B' is also likely to be larger than that for A'. Because both A and A' are credit worthy, the marginal difference in terms is likely to be small—they are both have access to as much credit as desired at relatively similar terms. The difference between B and B', however, could be the difference between access to credit and being rationed out of the market entirely.[129] Further, with diminishing marginal utility of wealth, the gains to B' from separation are likely to be more highly valued. Figure 6 illustrates this scenario in the context of the framework from Part III. The top panel is the prime market, and the bottom panel is the subprime. The losses to the prime market from lack of separation between A and A' are trivial, as reflected by the relatively flat $x^*(R)$ for the prime market, where R is credit worthiness and higher levels of x are better terms. For the subprime market, on the other hand, $x^*(R)$ is steep, reflecting the fact that separating good from bad risks in the subprime pool will lead to large gains, as discussed above. These gains from separation relative to the prime market

[129] *See* Einav, Jenkins & Levin, *supra* note 49.

are further amplified by the greater dispersion of risks in the subprime pool,[130] which are illustrated with a wider distribution of types.

Figure 6: Prime and Subprime Credit Markets

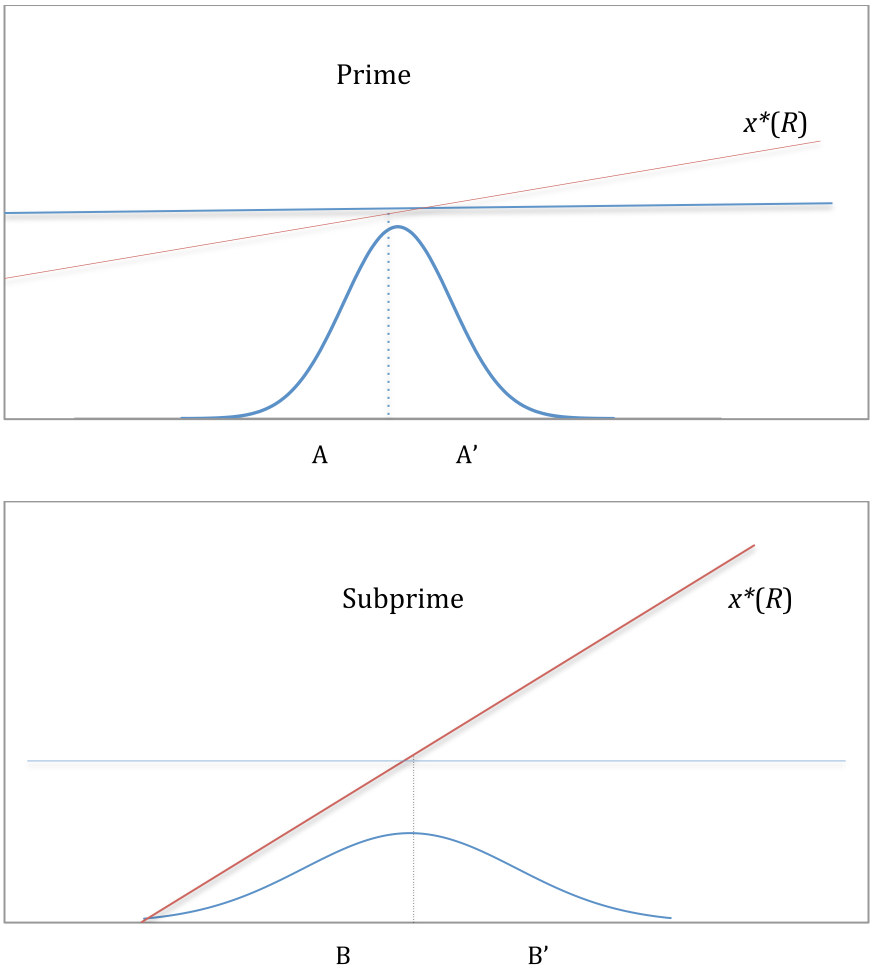

Empirical evidence on the impact of credit scoring and risk-based pricing on the poor's access to credit are suggestive of big data's potentially positive impact on those on the lower rungs of the economic ladder. For example, Federal Reserve data show that from 1983-2010, the largest increases in credit card ownership are in the bottom half of income earners (200-300%).[131] Moreover, from 1970-2010, there was a 77 increase in access to consumer credit by the lowest quintile compared to a 14 percent

---

[130] *See* Einav, Jenkins & Levin, *supra* note 49;Michelle A. Danis &Anthony Pennington-Cross, *The Delinquency of Subprime Mortgages*, 60 J. ECON. & BUS. 67, 74 (2008).
[131] *See* THOMAS A. DURKIN, GREGORY ELLIEHAUSEN, MICHAEL E. STATEN & TODD J. ZYWICKI, CONSUMER CREDIT AND THE AMERICAN ECONOMY, at 304 (2014).

increase for the highest quintile.[132] The poor also appear to gain from automated underwriting for mortgages as well. For example, a study finds that automated underwriting is more accurate at predicting risk than manual underwriting, and as a result approves more lower-income borrowers.[133] The authors conclude

> It is not surprising that the increased accuracy of [automated underwriting] benefits to a larger extent underserved populations. This group tends disproportionately to have higher-risk values for the attributes commonly used when underwriting mortgages. As a result, the poor stand to gain the most from AU's enhanced ability to better distinguish between low- and high-risk applicants of the margin of acceptable risk.[134]

A 2006 Federal Reserve Board report to Congress on credit scoring echoed these themes.[135] The study explained that credit scoring "could allow lenders to identify borrowers who are reasonable credit risks but who were previously underserved," and when coupled with risk-based pricing it had the potential to "expand the range of applicants to whom lenders are able to make loans profitably."[136] The data bore out these predictions. The report found that the credit use gap between low- and middle-income populations and high-income populations shrank from 1983-2004, and that in any event there was no evidence that those in the top of the income distribution disproportionately gained from increased information about credit history.[137]

A recent study examining automobile loans to subprime populations also illustrates how providing lenders with more granular information about borrowers can help lower-income populations by allowing them to identify relatively less risky borrowers within a risky population.[138] A used car seller dealt with very financially distressed population: average annual income was $28,000, and default rates on loans were over 60 percent. Further, there was strong evidence the data of both moral hazard and adverse selection:

---

[132] *Id*. at 72.
[133] *See* Susan Wharton Gates, Vanessa Gail Perry & Peter M. Zorn, *Automated Underwriting: Good News for the Underserved?* 13 HOUSING POLICY DEBATE 369 (2002).
[134] *Id.* at 385. A similar study finds that credit scoring and automated underwriting increased the amount of small business loans in low- and moderate-income census tracks by $.5 billion. W.S. Frame, Machael Padhi & Lynn Woosley, *Credit Scoring and the Availability of Small Business Credit in Low- and Moderate-Income Areas*, 39 FIN. REV. 35 (2004).
[135] BOARD OF GOVERNORS OF THE FEDERAL RESERVE SYSTEM, *REPORT TO THE CONGRESS ON CREDIT SCORING AND ITS EFFECTS ON THE AVAILABILITY AND AFFORDABILITY OF CREDIT* (Aug. 2007).
[136] *Id.*
[137] *Id.*
[138] Einav, Jenkins & Levin, *supra* note 49.

default rates increased substantially with loan amounts, and those who presented the greatest risk of default *ex ante* tended to ask for larger loans. Prior to credit scoring, this dealer offered only one flat rate and a capped loan amount. Once the dealer was able to use credit scores to more finely determine credit risks, however, it was able to extend more credit to those within this population who were relatively more credit worthy. By reducing defaults, moreover, the lender increased profits. Indeed, there was a large variation in default risk within this population, with the most risky borrowers about twenty percentage points more likely to default than the least risky.[139] The value of these data, according the authors was the ability to separate "consumers with transitory bad records from persistently bad risks."[140] These results should not be surprising. When adverse selection leads to credit rationing, those at the bottom rung of the ladder are the ones to suffer the most severe constraints on credit.

Another recent study on the inclusion of alternative data—e.g., utility and cell phone payment history—in credit scoring further suggests that to the extent that big data brings more information to bear on underserved populations, it is likely to benefit them.[141] Many lower-income consumers have little or no information on file with credit reporting agencies. As a result, lenders are unable to make reliable inferences about them: "unscorable" consumers typically viewed as high risk, and, so-called "thin file" consumers – those for whom there are few trade lines – are placed in lower credit tiers that they typically deserve.[142] Using credit files from three major credit reporting agencies, the authors find that inclusion of alternative data overwhelmingly increases the credit scores of thin file and unscorable consumers. Applying these credit score changes to estimate changes in access to credit, the authors find that lower income acceptance rates rise by 20 percent, compared to only a 5 percent increase for the highest income group.[143] The authors conclude:

> Members of lower income households benefit much more from the use of alternative data than members of higher income households. This is not surprising since it is the case that members of lower income households make up a disproportionately large share of the credit underserved, specifically those consumers with no credit files or thin credit files.[144]

---

[139] *Id.*
[140] *Id.* at 255. This is the difference between consumers B and B' in Figure 2.
[141] Michael A. Turner, Patrick D. Walker, Sukanya Chaudhuri & Robin Varghese, *A New Pathway to Financial Inclusion: Alternative Data, Credit Building, and Responsible Lending in the Wake of the Great Recession*, POLITICAL & ECONOMIC RESEARCH COUNCIL (2012).
[142] *See Id.* at 13.
[143] *Id.* at 17.
[144] *Id.*

Along these lines, several start-ups are using big data to analyze thousands of variables like rent records, prior payday loans, pawnshop transactions, Facebook friends and other Internet footprints to identify better credit risks within poor populations.[145] For example, Zest—a company started by Google's former chief information officer—claims that by using big data to analyze records sourced from individuals' social network and internet footprints, those who have traditionally been denied credit due to lack of information about them in the system can see their credit scores rise by up to 40 percent.[146] The upshot is that these alternative scoring systems can give underserved populations alternatives to payday lenders or pawnshops, and by identifying creditworthy individuals, lenders are dramatically reducing default rates below those experienced by payday lenders.[147]

Just as the benefits to the poor stand to be large, the intrinsic privacy concerns appear to be minimal. The type of information feeding the algorithms is relatively non-sensitive, and the prediction—creditworthiness—has primarily strategic value in concealment. That is, only those with poor scores would want to conceal this information. Consequently, the use of big data here appears to present a circumstance analogous to distribution f_1 and z_H from Figure 3, and suggests that hands-off regulatory approach.

B *Price Discrimination*

Next, consider concern voiced by several authors that big data-driven price discrimination will cause the poor to pay more than the rich.[148] Indeed, a consistent theme in the existing privacy scholarship is that firms armed with Big-Data will be able to extract ever-increasing amounts of consumer surplus.[149] Price discrimination comes in three varieties, conveniently labeled first-, second-, and third-degree. First-degree discrimination is often referred to as "perfect" price discrimination, as it involves a firm charging each consumer his or her exact willingness to pay. While this type of discrimination leaves inframarginal consumers worse off,

---

[145] *See* Elizabeth Dwoskin, *'Big Data' Doesn't Yield Better Loans*, WALL ST. J, Mar. 17 2014, http://www.wsj.com/articles/SB10001424052702304732804579425631517880424.
[146] *See* Patrick Jenkins, *Big Data lends new Zest to banks' Credit judgment*, FIN. TIMES, June. 23, 2014; Quentin Hardy, *Big Data For the Poor*, N.Y. TIMES, July 5, 2012.
[147] *Crunching the Numbers*, ECONOMIST at 7 (May 19, 2012).
[148] *See* Tene, *supra* note 32 (discounts to the rich subsidized by price hikes for the poor); Angwin & Soltani, *Websites Vary Deals and Prices Based on Users' Information*, WALL ST. J, (Dec. 12, 2012) (finding that differential online pricing based on zip code leads to those in relatively poorer zip codes to pay more).
[149] *See, e.g.*, Calo, *supra* note 9, at 33 (firms will use big data to charge consumers "as much as possible" and to manipulate them to buy products and services that they "[do] not need or need[] less of."); Peppet, *supra* note 21 (we know that poor rarely win from competition).

it unambiguously increases welfare because it expands output; consumers whose willingness to pay falls below the uniform price, but above the marginal cost of production, were previously priced out of the market and now are able to participate at lower prices.[150]

Because of data demands, first-degree price discrimination is mostly relegated to domain of theory. Firms instead rely chiefly on less fine market segmentations, either by allowing consumers to self-select based on non-linear pricing schemes or product attributes (second-degree), or by using observable characteristics like age as proxies for willingness to pay to segment markets (third-degree). Although a detailed treatment of the welfare effects of second and third-degree discrimination is well beyond the scope of this paper, suffice to say, it's complicated.[151] A necessary condition for price discrimination to be welfare-enhancing is that it spur an increase in output, a condition that fits neatly into the framework presented in Part II. If R is willingness to pay, assembling information to identify or predict R is efficient only if it prompts firms to make sufficiently low offers to draw marginal consumers into the market (x), which increases $V(x)$. If instead $\frac{\partial x*}{\partial R} = 0$, or the cost of identifying R is greater than the additional value created by bringing marginal consumers into the market, price discrimination—and the use of big data to identify consumer valuations—is dissipative, as firms would be spending c merely to transfer t from consumer to themselves.

Although the welfare effects of second- and third-degree price discrimination are indeterminate theoretically, there is a widespread view among economists and antitrust enforcers that price discrimination is at worst benign and probably welfare-enhancing; a view that is bolstered by empirical evidence.[152] Neither federal antitrust agency has enforced the Robinson-Patman Act in decades,[153] and the Department of Justice sided with the defendant in the most recent Robinson-Patman case heard by the Supreme Court, arguing that a ban on price discrimination was likely to harm competition.[154] Further, the bi-partisan Antitrust Modernization

---

[150] First-degree price discrimination can be welfare-reducing if the discriminating firm invests more in effecting discrimination than is gained from reduction in dead weight loss.

[151] *See* HAL R. VARIAN, MICROECONOMIC ANALYSIS 250-53 (3d ed. 1992).

[152] *See, e.g.,* Igal Hendel & Aviv Nevo, *Intertemporal Price Discrimination in Storable Goods Markets*, 103 AM. ECON. REV. 2722 (2013); P. Leslie, *Price Discrimination in Broadway Theory*, 35 RAND J. Econ. 520 (2004); Andrew Cohen, *Package Size and Price Discrimination in the Paper Towel Market*, 26 INT'L J. INDUS. ORG. 502 (2008).

[153] The Antitrust Division has not brought a Robinson-Patman case since the 1960s, and the FTC has brought only one Robinson-Patman case since 1992. *See* ANTITRUST MODERNIZATION COMMISSION, REPORT & RECOMMENDATIONS at 318 (2007).

[154] *See* Brief for the United States as Amicus Curiae Supporting Petitioner at 27 & n.15, *Volvo Trucks N. Am., Inc. v. Reeder-Simco GMC, Inc.*, 544 U.S. 164 (2006) ("Imposing liability for differences in concessions offered to dealers bidding on different sales would limit suppliers'

Commission recommended the repeal of the Robinson-Patman Act, concluding:

> [S]eventy years after passage of the Robinson-Patman Act, courts remain unable to reconcile the Act with the basic purpose of antitrust laws to protect competition and consumer welfare. . . There is no point in further efforts to reconcile the Act with the antitrust laws in general; the Robinson-Patman Act instead should be repealed.[155]

It is also important to note, is that as we move from a world in which firms rely on crude proxies for willingness to pay—age, income, purchase of complementary goods etc.—towards more granular targeted pricing, we begin to move toward a world of first-degree price discrimination, which is unambiguously welfare-enhancing.[156] Importantly, because income is negatively related to willingness to pay, the poor are exactly the ones who are most likely to gain as price discrimination becomes easier to implement. Assertions that price discrimination brought about by big data is likely to allow firms to implement schemes under which the poor to subsidize the rich are just poor economics. If a firm can segment markets, optimal pricing requires the market with the most elastic demand to pay the lower prices.[157] Because price elasticity of demand is a negative function of income, a firm that segments its market into rich and poor consumers would charge a higher price to the former and lower one to the latter;[158] think student or elderly discounts at movies and restaurants, or the Saturday stay-over and advance booking requirements for cheaper flights.[159] Indeed, one of the few attempts to use big data to price discriminate that became public involved Orbitz placing higher-priced hotels more prominently in search results for

---

ability to tailor prices to the competitive situation, and thus diminish the vigor of interbrand price competition.").

155 ANTITRUST MODERNIZATION COMMISSION REPORT, *supra* note 129, at 322.

156 This effect is analogous to that recognized by Strahilevitz in conjunction with statistical discrimination. Lior Jacob Strahilevitz, *Privacy versus Antidiscrimination*, 75 U. CHI. L. REV. 363 (2008). Strahelivitz argues that as we move from a world in which parties use protected classes as crude proxies for undesirable economic characteristics to one in which they can measure undesirable economic characteristics directly, statistical discrimination is likely to decline.

157 This is called Ramsey pricing, and formally requires: $\frac{P_A}{P_B} = \frac{1+\frac{1}{\varepsilon_A}}{1+\frac{1}{\varepsilon_B}}$, where ε_i is the own-price elasticity of demand for good *i*. DIETER BÖS, PRICING AND PRICE REGULATION: AN ECONOMIC THEORY FOR PUBLIC ENTERPRISES AND PUBLIC UTILITIES (3d ed. 1994).

158 Studies show, for example, that the poor respond to excise taxes on cigarettes and alcohol by curtailing their consumption more than the rich. Michael Grossman, Frank J. Chaloupka & Richard Anderson, *A Survey of Economic Models of Addictive Behavior*, 28 J. OF DRUG ISSUES 631, 635 (1998).

159 *See* N. Gregory Mankiw, Principles of Microeconomics (Joseph Sabatino et al. eds., 6th ed. 2011).

Mac users under the assumption that Mac users typically are wealthier than PC users.[160]

The discussion above limited the analysis to one firm's pricing in isolation. But firms' actions do not take place in a vacuum; competition is all but ignored in the standard treatment of big data's impact on consumers. Although firms rationally seek to extract as much surplus as they can from consumers, they are limited in this quest by the fact that in most markets several other firms are trying to accomplish the same thing.

For example, consider the following example to see how interjecting competition into the standard big data-driven price discrimination story dramatically alters its conclusions. Figure 7 shows a Hotelling line with Lands End on one end and L.L. Bean on the other. Consumers are arrayed along the line (with length of one), with those near the left having the strong preferences for L.L. Bean's clothes, and those near the right end having a strong preference for Lands End clothing. Consumers near the middle are largely indifferent between the two stores. Suppose that each seller's marginal cost for an oxford shirt is $10, that consumers value their ideal oxford shirt at $25, and that they suffer $10 in disutility for each unit they have to consume away from their position on the line. It can be shown that the equilibrium price for a shirt will be $20, which is determined by L.L. Bean and Lands End competing for the marginal consumers in the middle.[161]

Figure 7: Spatial Competition

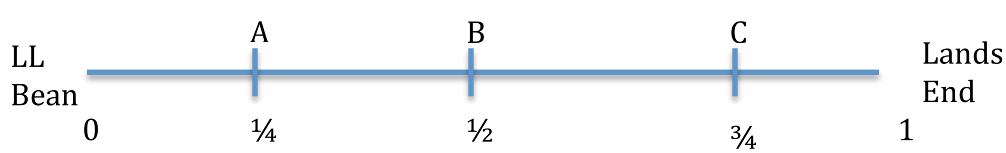

---

[160] This instance was not really price discrimination because the Mac users were charged the same prices as PC users for the same hotel. More expensive hotels were just more prominently placed for the Mac users. Dana Mattioli, *On Orbitz, Mac Users Steered to Pricier Hotels*, WALL ST. J. (AUG, 23, 2012), HTTP://WWW.WSJ.COM/ARTICLES/SB10001424052702304458604577488822667325882.

[161] This equilibrium is derived by assuming Bertrand competition between L.L. Bean and Lands End over consumers with utility functions:
$U_{LL} = \$25 - 10\tau - P_{LL}$ and $U_{LE} = \$25 - 10(1-\tau) - P_{LE}$, where subscripts LL and LE are utility and price associated with purchasing from L.L. Bean and Lands End, respectively, and τ is the distance of a consumer's ideal point from L.L. Bean in product space.

Suppose now that big data allows these firms to peer into consumers' minds and understand exactly where they resided along the line. First, focus on L.L. Bean's decision with respect to consumers A, B, and C located at positions ¼, ½, and ¾, respectively. Bean now knows that it can charge A the most ($22.50), because he has a strong preference for the Bean brand. By the same token, Bean knows that C has relatively weak preferences for Bean, but can be lured with a price of $17.50. B is the marginal consumer, and will receive the same price as he did in the uniform price equilibrium, $20. This is typically where the big data price discrimination story ends – those with higher values suffer higher prices, and even those with lower values who are brought into the market have their entire surplus extracted.

In most markets, however, this isn't where the story ends. It's unlikely that Lands End will sit idly by as L.L. Bean poaches its customers. Further, armed with the same information, Lands End rationally will try to expand its market to compete for L.L. Bean's customers. In the end, Lands End and L.L. Bean compete for all customers along the line, and as a result everyone pays lower prices than they would in a uniform-price equilibrium. The marginal consumers at B pay $10, and consumers at points A and C each pay $15 as Bean and Lands End are forced to compete for previously captive consumers. That is, when competition is considered, the ability to target prices can increase consumer welfare.[162] The extent to which this type of targeting currently occurs is unclear, but it has been seen in the grocery store for years. If you buy Dannon yogurt, for example, it is not be uncommon for your receipt to include a coupon for Yoplait or another competing brand. Similarly, it has become common for firms to bid on rival trademarks as keywords in search advertising. For example, Lens.Com may bid for the term "1-800" so that its ads appear to consumers searching for 1-800 Contacts.[163] All of this suggests that hypotheticals involving the use of big data to target vulnerability individuals (e.g., coupons for donuts to those trying to diet) are incomplete;[164] they need to be reworked to account for strategic behavior by competitors.

---

[162] This equilibrium arises because the distant seller will set its price equal to marginal cost for all consumers that are outside of its market. The best response to this pricing strategy for L.L. Bean is $p = 10(1 - 2\tau) + 10$, and $p = 10(2\tau - 1) + 10$ for Land End. In this simple model, although consumers are unambiguously better off, total welfare remains unchanged because output remains unchanged. More general models show that total welfare can increase when markets exhibit "best response asymmetry" – i.e., one firm's weak market is another's strong market – and they place relatively more weight on their strong market. *See* Kenneth S. Corts, *Third Degree Price Discrimination in Oligopoly: All-Out Competition and Strategic Commitment*, 29 RAND J. ECON. 306 (1998); Lars A. Stole, *Price Discrimination & Competition* in 3 HANDBOOK OF INDUSTRIAL ORGANIZATION (2007); Thisse & Vives, *On the Strategic Choice of Spatial Price Policy*, 78 AM. ECON. REV. 122 (1998).

[163] David A. Hyman & David J. Franklyn, *Trademarks as Search Engine Keywords: Who, What, When?*, 92 TEX. L. REV. 2117 (2014); *1-800 Contacts, Inc. vs. Lens.com, Inc.*, 2013 WL 3665627 (10th Cir. 2013).

[164] *See, e.g.,* Calo, *supra* note 9, at 1031-34.

Leaving aside differential pricing, big data can enhance competition merely by allowing firms to make themselves visible to customers. Imagine a world in which L.L. Bean was the established incumbent and few knew of Lands End. Even if consumers do not know about Land End, Lands End can use data to find consumers who may like their clothing. In this manner, big data-driven algorithms that predict consumer tastes are likely to enhance competition by reducing the costs to consumers of finding competitors. There is a large literature on search costs that suggest when consumers can become aware of competing offers more easily, average prices tend to fall as does price dispersion.[165] Further, the role that credit scoring has played in promoting competition in consumer credit markets may be instructive. Prior to the widespread adoption of credit scoring, most consumers had limited options for credit. Typically, consumers could chose only from local institutions with which they had a relationship. Once credit reporting became widespread, national institutions could target consumers all over the country.[166] In response, local lenders were forced to reduce their rates. What the consumer saw was lower rates and annual fees, and an explosive increase in the availability of credit.[167] Just as credit scoring and risk-based pricing expanded consumer options, predictive analytics made possible by big data are likely to expand the ability of firms of to reach new consumers.

In the context of the framework from Part III, there is sufficient reason to believe that big data-driven price discrimination is likely to increase welfare ($\frac{\partial x*}{\partial R} > 0$), and that the gains may accrue disproportionately to the poor. Further, the data involved—both input and output—are unlikely to raise serious intrinsic privacy costs. Accordingly, like credit markets, price discrimination also appears to present a circumstance analogous to distribution f_1 in Figure 3. And even if benefits are represented by z_L, a hands-off regulatory approach is still suggested.

C. *Labor Markets*

---

[165] *See* Steven Salop & Joseph Stiglitz, *Bargains and Ripoffs: A Model of Monopolistically Competitive Price Dispersion*, 44 REV. ECON. STUD. 493 (1977); Dale O. Stahl II, *Oligopolistic Pricing with Sequential Consumer Search*, 79 AM. ECON. REV. 700 (1989); Kenneth Burdett & Kenneth L. Judd, *Equilibrium Price Dispersion*, 51 ECONOMETRIC SOC. 955 (1983); Maria Arbatskya, *Ordered Search*, 38 RAND J. ECON. 119 (2007).

[166] *See* Durkin et al., *supra* note 107, at 268-69.

[167] *See Id.* at 270. Knittel & Stango find evidence that the new entry made available by credit scoring help break a pattern of tacit collusion among banks that led to price stability for over a decade. Christopher R. Knittel & Victor Stango, *Price Ceilings as Focal Points for Tacit Collusion: Evidence from Credit Cards*, 93 AM. ECON. REV. 1703 (2003).

Finally, big data also has the potential to ameliorate socially wasteful information asymmetries in job markets. Again, these benefits largely may accrue to those at the lower-end of the socioeconomic scale who have been excluded from labor market gains that have gone predominantly to those with post-secondary education.[168] Many entry level jobs require post-secondary education that is unrelated to the skills the job requires, suggesting that educational investments are a signaling mechanism that helps employers sort candidates into high- and low-productivity bins. To the extent that individuals are required to make larger investments in education than they otherwise would, it represents a social waste. As discussed in Part II, a large body of empirical work lends support to the signaling value of education.[169]

Some companies are beginning to use big data analytics to identify candidate employees for tech, high-end sales, and managerial positions, and these analytics are suggesting that other indicators are more predictive of good fits than college.[170] In this manner, big data can allow those without a post-secondary education to compete for jobs that previously were open only to college graduates. Here, R is expected productivity and $x(R)$ is an offer that matches the type. To the extent that such uses of big data improve labor market matching, $\frac{\partial x*}{\partial R} > 0$, which implies that as long as the costs are not too high, this kind of data use is socially productive. Reductions in expenditures on post-secondary education intended primarily to serve as a signal of productivity, moreover, would increase welfare for everyone by allowing these resources to be put to more productive use.

The benefit of big data in this circumstance, again, likely would accrue primarily to those on the lower rungs of the economic ladder—those who could not afford post-secondary education. The data are clear that the largest share of economic gains over the past three decades have gone to those with college degrees: the ratio of mean earnings for college grads to high school grads has risen by 31 percent for men and 39 percent for women since 1980.[171] To the extent that big data can unlock the doors to jobs that that were previously only for those with a college education, it may allow a wider sharing of economic gains.

---

[168] *See* Josh Zumbrun, *Just How Stagnant are Wages Anyway?*, WALL STREET JOURNAL (Jul. 6, 2015), at http://blogs.wsj.com/economics/2015/07/06/just-how-stagnant-are-wages-anyway/.
[169] *See* note 59 *supra*, and accompanying text.
[170] *See* Don Peck, *They're Watching You at Work*, THE ATLANTIC (Dec. 2013); notes 68-70, *supra* and accompanying text.
[171] *See* Bureau of Labor Statistics, Changes in the College/High School Earnings Differential, by Gender, 1970 – 2015, *at* www.bls.gov.

V. CONCLUSION

Calls to regulate big data must be careful to distinguish between productive and dissipative concealment of private facts. Retarding the use of data to draw inferences merely because they will cause some people to suffer worse terms is unambiguously welfare-reducing, and it's unfair: it forces those with relatively good attributes to subsidize both those with bad attributes and extreme preferences for privacy. What's more, claims that big data is likely to harm the poor do not stand up to close scrutiny. Preventing discovery of relatively good types in disadvantaged populations only condemns them to suffer higher prices, less credit, and fewer job opportunities. Given the lack of knowledge about how consumers suffer intrinsic privacy harms, policy generally should tread lightly, with a focus on use rather than collection. Big data raises genuine privacy concerns, but anxieties over so-called predictive privacy harms risk confusing benefits for costs.