

VIRGINIA JOURNAL OF LAW & TECHNOLOGY

FALL 2017

UNIVERSITY OF VIRGINIA

VOL. 21, No. 01

Separation Anxiety

JAMES C. COOPER¹

¹ Associate Professor of Law and Director, Program on Economics & Privacy, Antonin Scalia Law School at George Mason University. I thank Jane Bambauer, Eric Claeys, Bruce Kobayashi, Craig Lerner, Koren Wong-Ervin, Todd Zywicki, and participants at the George Mason University School of Law Levy Workshop for helpful comments on earlier drafts. I also thank John Magruder for outstanding research assistance and Roger Gibboni for valuable editorial assistance.

ABSTRACT

Privacy is about being “let alone,” so in one sense, privacy means to separate yourself from the world. Paradoxically, by concealing facts about yourself, observers view you as less separated from everyone else. They can no longer make out the features that distinguish you from those to whom you bear a superficial resemblance. In this manner, privacy promotes pooling. Markets, however, tend to benefit from separation—the ability to distinguish between different types. This tension between pooling and separation is on full display in the privacy law scholarship surrounding big data, which tends to see separation as a cause for concern. Privacy is valuable. It’s woven into the fabric of our society. But we must be careful to discern between privacy’s intrinsic and strategic values before prescribing drastic measures to address this separation anxiety. Strategic privacy concerns redistribution and merely destroys surplus. Intrinsic privacy, on the other hand, satisfies a demand and hence increases welfare. The major contribution of this Article is to develop a positive framework based on the economics of contracts and torts to identify when limiting separation may be justified. I find that when strategic privacy is at issue, policy should be rooted in antidiscrimination law—which embodies the choices that society has made about which traits are fair game for classification—rather than privacy law. Alternatively, privacy law should be used when intrinsic privacy is implicated. The analysis suggests that *ex ante* restrictions on use make sense only in the narrow circumstances in which there is likely to be agreement that the big data predictions implicate highly sensitive information. Alternatively, when there is little agreement on how privacy harms are likely to be suffered, the default regulatory posture should be one of notice of collection and use, with the Federal Trade Commission enforcing a firm’s promises.

TABLE OF CONTENTS

I. INTRODUCTION	3
II. THE PROBLEM.....	9
III. ASYMMETRIC INFORMATION: ADVERSE SELECTION & MORAL HAZARD	16
A. Separation, Pooling, and Adverse Selection	17
B. Dynamic Considerations: Moral Hazard and Endogenous Types	20
IV. BIG DATA	22
A. Adverse Selection	23
B. Moral Hazard	24
C. Benefits for the Poor.....	25
V. SEPARATING BENEFICIAL FROM HARMFUL USES OF BIG DATA	31
A. Framework.....	31
B. Dissipative Privacy and Dissipative Revelation	36
1. Dissipative Privacy	38
2. Dissipative Privacy and Antidiscrimination	40
3. Dissipative Revelation	41
C. Concealment and Revelation Both Valuable.....	42
1. Factors Influencing Gains from Separation.....	42
2. Identifying Optimal Restrictions on Big Data.....	44
3. Measuring Intrinsic Privacy Harm.....	47
4. Developing Defaults.....	50
i. Default in Favor of Big Data Classification	51
ii. Default Against Big Data Classifications.....	53
iii. The Hard Cases	55
VI. CONCLUSION.....	57

INTRODUCTION

Privacy is about being “let alone.”² So in one sense, privacy means to separate yourself from the world. But perhaps paradoxically, by concealing facts about yourself, observers view you as *less* separated from the rest of the world. They can no longer make out the features that differentiate you from others to whom you bear a superficial resemblance. In this manner, privacy promotes “pooling”—when facts that distinguish you from the rest of the world are obscured, you appear to be in the same pool as everyone else.

Markets, however, tend to benefit from “separation”—the ability to distinguish between different types. Imagine the automobile insurer who is unable to query driving records when writing policies. The inability to separate good from bad drivers means that all drivers are pooled together and pay a single rate based on average risk. It’s clear here that good drivers subsidize the risky ones, but that is not the end of the story. It’s not a zero-sum game. Some good drivers will exit the market, as the rates are too high. In the end, only the riskiest drivers remain.³ This phenomenon is called “adverse selection”—too many risky types, and too few good types, select into the market. Costs go up, the number of insured goes down, and society is worse off as a result. Pooling is wasteful in other ways as well. Insulating risky types from the impact of their decisions creates “moral hazard” by dulling incentives to invest in becoming a “good” type.⁴ For example, a smoker who has to reveal his habit will pay higher rates than he would if he could pool with non-smokers and therefore may have a greater incentive to quit. Further, when there is pooling, good types devote resources to try to convince others that they really are good, and those in the dark invest in mechanisms to determine who is telling the truth when they declare that they are good and who is lying.

² Samuel Warren & Louis Brandeis, *The Right to Privacy*, 4 HARV. L. REV. 193, 193 (1890).

³ See George Akerlof, *The Market for Lemons: Quality Uncertainty and the Market Mechanism*, 84 Q. J. ECON. 488 (1970).

⁴ 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 for those at the bottom of the economic rung who are relatively better risks than the rest of their cohort. See *infra* Parts III.A-B and accompanying text.

Economists have won several Nobel Prizes by studying these problems associated with informational asymmetries.⁵

This tension between market efficiency and privacy—between separation and pooling—is on full display in the privacy law scholarship surrounding big data, the use of the ever-growing data stream to make predictions about us.⁶ 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, it has 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.⁷ In this manner, separation itself is a cause for anxiety.

What tends to emerge from this literature is increasing calls to limit firms’ ability to improve inferences about those with whom they deal. For example, some have suggested limiting consumers’ ability to acquiesce to monitoring to prevent some from suffering the negative inference that a firm could draw from an unwillingness to be monitored.⁸ Others have suggested “due process” rights in big data predictions, likening these determinations to government deprivations of liberty.⁹ It is not really an exaggeration to say that this line of legal scholarship views separation as a harm and pooling as the remedy.

⁵ E.g., George Akerlof, Joseph Stiglitz, Michael Spence.

⁶ For some of the earliest expressions of the tension between market efficiency and privacy, see George J. Stigler, *An Introduction to Privacy in Economics and Politics*, 9 J. LEGAL STUD. 623 (1980); Richard A. Posner, *The Right of Privacy*, 12 GA. L. REV. 393 (1977).

⁷ 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, Fed. Trade Comm’n, Big Data and Consumer Trust: Progress and Continuing Challenges, Remarks Before the International Conference of Data Protection and Privacy Commissioners (Oct. 15, 2014).

⁸ Peppet, *supra* note 7, at 1158–59.

⁹ Crawford & Shultz, *supra* note 7, at 118.

This thinking, moreover, has also begun to seep into policy discussions. The White House, for example, recently floated a draft privacy bill, which would establish “privacy review boards” to clear big data analysis that has the potential to result in “adverse actions concerning multiple individuals.”¹⁰

At first glance, this separation anxiety stands what most economists are taught in graduate school on its head—as a general matter, policies that force pooling squander valuable information.¹¹ Of course, merely to say that privacy retards information flows is insufficient to condemn it. Privacy without a doubt is valuable. It’s woven into the fabric of our society—privacy rights are embedded in our Constitution¹² and the common law.¹³ The debate sparked by the San Bernardino shooter’s iPhone illustrates how important privacy is to U.S. consumers.¹⁴ But in addressing this separation anxiety, we must be careful to discern exactly what gives rise to privacy’s value before prescribing drastic *ex ante* restrictions on information flows.

More specifically, we must be careful to distinguish between privacy’s *intrinsic* and *strategic* values. Intrinsic value refers to the direct utility one derives from being free from unwanted observation or being able to limit the knowledge of certain personal facts to oneself or a close circle of friends and family. Strategic privacy, on the other hand, is wholly dissipative. It is the value that accrues to a party from obfuscating facts relevant to a transaction in hopes of getting a bigger

¹⁰ *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>.

¹¹ There are exceptions. When the distributional gains from separation are greater than the productive gains, resources spent on separation—either signaling or screening—are socially wasteful. See *infra* note 128 and accompanying text.

¹² See e.g., *Griswold v. Connecticut*, 318 U.S. 479 (1965); *Whalen v. Roe*, 429 U.S. 589 (1977); see also DANIEL J. SOLOVE & PAUL M. SCHWARTZ, *INFORMATION PRIVACY LAW* 35–36 (5th Ed. 2015) (discussing privacy elements in 1st, 3rd, 4th, and 5th Amendments).

¹³ Most states have some form of privacy torts, which include false light, publication of private fact, intrusion into seclusion, and appropriation of name or likeness. See *RESTATEMENT (SECOND) OF TORTS* § 652 (1977).

¹⁴ See Devlin Barrett, *Americans Divided Over Apple’s Phone Privacy Fight*, *WSJ/NBC Poll Shows*, *WALL ST J.*, Mar. 9, 2016, available at <http://www.wsj.com/articles/americans-divided-over-apples-phone-privacy-fight-wsj-nbc-poll-shows-1457499601>.

slice of the pie.¹⁵ The prospective borrower who conceals his plans to quit his job is more likely to get a lower interest rate. The prospective employee who hides the fact of her drug addiction is more likely to get the job. It's important to emphasize that strategic privacy does not merely transfer value from lender to borrower or employer to employee; it exacerbates the problems discussed above associated with adverse selection and moral hazard. Put differently, when privacy serves purely strategic ends, losses to risky types due to big data-driven sorting should never be counted as privacy harm because they are merely artifacts of a reduction in asymmetric information. Without these losses, the net gains to society cannot materialize.

One observation that flows readily from the distinction between strategic and intrinsic privacy is that the former really doesn't concern privacy at all. Accordingly, if we want to prevent classification based on traits that big data can ferret out, the mechanism should be rooted in antidiscrimination law¹⁶—which embodies the choices that society has made about which traits are fair game for classification—rather than privacy law.¹⁷ Indeed, an artifact of being forced to think through the traits on which society will forbid separation is that we map out its complement—the set of traits that we permit big data to discover because their concealment provides no social value. That is, when privacy is solely strategic, unless there are other socially beneficial reasons for preventing discrimination on the trait in question, it makes no sense to bar such classification on privacy grounds. For example, although lenders cannot base terms on race, gender, or religion, it would be odd—and socially wasteful—to prevent them from discriminating based on credit history. Using antidiscrimination rather than privacy to force pooling also has the advantage of discouraging wasteful investments in signaling. Even if privacy law prevents a firm from attempting to find the value of some hidden trait, nothing prevents a consumer from attempting to signal her value of that trait.¹⁸ However, if a firm is forbidden from making decisions (*e.g.*, hiring, pricing, insurance) based on the trait in question, signaling no longer has value; if firms cannot make hiring

¹⁵ See, *e.g.*, Posner, *supra* note 6.

¹⁶ See, *e.g.*, Equal Credit Opportunity Act, 15 U.S.C. § 1691 (2014); Fair Housing Act, 42 U.S.C. § 12101 (2009); Genetic Information Nondiscrimination Act, 42 U.S.C. § 2000ff (2008).

¹⁷ *E.g.*, The Federal Trade Commission Act, 15 U.S.C. § 45 (2006).

¹⁸ See Peppet, *supra* note 7, at 1156.

decisions based on health, there's no reason to join the gym before your interview. Further, an antidiscrimination regime would save resources that risky types would devote to concealing negative values of the trait in question, which would increase the amount of information available to society.

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, however, it may be difficult to disentangle the inherent and strategic dimensions of privacy. For example, health privacy has a strategic dimension: keeping the fact of diabetes or bipolar disorder secret likely will lead to advantages in the labor and insurance markets. But most people also derive intrinsic value from keeping health conditions private. In these cases, where there is a legitimate privacy interest at stake, it is appropriate for policy makers to use privacy laws to regulate big data. However, the thin extant literature on valuation of privacy provides little guidance on the likely privacy benefits of policies that retard big data.¹⁹ What's more, the inherent valuation of privacy will vary across individuals and within individuals across contexts. 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.²⁰ Some may,²¹ but there is a distribution, and we do not know it. Nonetheless, this episode serves as a rallying cry for big data regulation.²² A more careful weighing of the costs and benefits of this use of big data is needed before we condemn such practices as archetypes of big data gone bad.

¹⁹ See *infra* notes 160-167.

²⁰ See Kashmin Hill, *How Target Figured Out A Teen Girl Was Pregnant Before Her Father Did*, FORBES, Feb. 16, 2012, available at <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, available at 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.

²¹ 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, http://www.salon.com/2014/04/28/one_womans_attempt_to_hide_her_pregnancy_from_big_data/.

²² See, e.g., Ryan Calo, *Digital Market Manipulation*, 82 GEO. WASH. L. REV. 995 (2014).

The major contribution of this Article is to provide a positive framework for thinking about big data regulation in these circumstances—when intrinsic and strategic privacy values are mixed and there are benefits to separation. 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. I use this framework to identify circumstances in which privacy and revelation are socially wasteful and also to suggest optimal policy defaults when big data reduces privacy that has both strategic and intrinsic dimensions. The analysis suggests that *ex ante* restrictions on use make sense only in the narrow circumstances in which there is likely to be agreement that the big data predictions implicate highly sensitive information. Alternatively, when there is little agreement on how privacy harms are likely to be suffered, the default regulatory posture should be one of notice, with the Federal Trade Commission enforcing a firm’s promises. Along the way, I also attempt to allay some of the fears 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 hypotheticals.

The remainder of this Article 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. Part III explores how big data is likely to ameliorate some of these problems. It also addresses claims that big data is likely to have a disproportionately negative impact on the poor and shows that there is reason to believe that the opposite is true. Drawing on the distinction between strategic and intrinsic values of privacy, Part IV constructs a framework for analyzing privacy harms from big data and suggests factors that are likely to militate for or against various regulatory defaults. The final part summarizes the main points and offers some policy prescriptions.

I. THE PROBLEM

Much has been written on the topic, so I will only briefly describe big data to lay the groundwork for the remainder of the Article. Big data is a general catchall term for the analysis of enormous datasets—sets that may even satisfy the condition that “ $N = \text{all}$ ”²³—to tease out correlations and relationships that could not be seen with small data sets.²⁴ 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.²⁵ 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 Flu Trends, an algorithm that predicts flu outbreaks based on Google search terms, is often held out as the quintessential example of big data.²⁶ Other notable examples include Oren Etzioni’s Farecast flight price prediction website,²⁷ Amazon’s book recommendation algorithm, Google Translate,²⁸ and Netflix’s movie

²³ VIKTOR MAYER-SCHÖNBERGER & KENNETH CUKIER, *BIG DATA: A REVOLUTION THAT WILL TRANSFORM HOW WE LIVE, WORK, AND THINK* 26 (2013); *see also* VIKTOR MAYER-SCHÖNBERGER & KENNETH CUKIER, *supra* note 23, at 6 (explaining that 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 . . .”).

²⁴ *Id.* at 12 (explaining that big data “is about applying math to huge quantities of data in order to infer probabilities”).

²⁵ STEPHEN DOYLE, *ESSENTIAL ICT A LEVEL: AS STUDENT BOOK FOR AQA 131* (2008).

²⁶ *See* GOOGLE FLU TRENDS – UNITED STATES, <https://www.google.org/flutrends/about/data/flu/us/data.txt>; *But see* Paul Ohm, *The Underwhelming Benefits of Big Data*, 161 U. PA. L. REV. ONLINE 339 (2013) (responding to Paul M. Schwartz, *Information Privacy in the Cloud*, 161 U. Pa. L. Rev. 1623 (2013)).

²⁷ Farecast was purchased by Microsoft and integrated into Bing but recently shut down. *See Farewell Farecast: Microsoft Kills Airfare Price Predictor, to the Dismay of its Creator*, *Geekwire*, GEEKWIRE (Apr. 18, 2014), <http://www.geekwire.com/2014/farewell-farecast-microsoft-kills-airfare-price-predictor-dismay-creator/>.

²⁸ Tim Harford, *Big Data: Are We Making a Big Mistake?*, *FINANCIAL TIMES* (Mar. 28, 2014), <http://www.ft.com/intl/cms/s/2> (“Google Translate is as close to theory-free, data-driven algorithmic black box as we have”).

recommendation algorithm.²⁹ Credit card companies also use big data methods to detect fraud by examining anomalies in purchasing patterns.³⁰ As will be discussed in more detail below, big data is 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.³¹

A great deal has been written attempting to classify harms that stem from data breaches or the unwanted collection and use of personal data.³² 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 is harmful. Then there are subjective harms, which include any direct psychic or embarrassment costs posed by online tracking of health care information or unwanted surveillance of intimate activities.³³ Because they are not objectively verifiable like monetary harms, and are likely suffered differently across populations and contexts, they are difficult to quantify. Nonetheless, they are harms in the traditional sense.³⁴

As big data has become the major focal point for 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

²⁹ See VIKTOR MAYER-SCHÖNBERGER & KENNETH CUKIER, *supra* note 23, at 110.

³⁰ See *Big Data: Crunching the Numbers*, ECONOMIST (May 19, 2012), <http://www.economist.com/node/21554743c>.

³¹ See *infra* Part IV.B and accompanying text.

³² See, e.g., Daniel J. Solove, *A Taxonomy of Privacy*, 154 U. PA. L. REV. 477 (2006); M. Ryan Calo, *The Boundaries of Privacy Harm*, 86 IND. L.J. 1131 (2012).

³³ See, e.g., Complaint, In re Designerware, LLC, Docket No. C4390 (F.T.C. April 15, 2013), available at <https://www.ftc.gov/sites/default/files/documents/cases/2013/04/130415designerwarecmt.pdf>.

³⁴ For example, the Supreme Court recently explained that intangible harms can form the basis for standing in federal court. See *Spokeo, Inc. v. Robins*, 136 S. Ct. 1540 (2016). Courts have read *Spokeo* to allow intangible harm from invasion of privacy interests. See, e.g., In re Vizio, Inc., Consumer Privacy Litigation, 2017 WL 1836366 (C.D. Cal. 2017).

make easy targets for hackers,³⁵ 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 speak, firms will be able to tailor offers based on your reconstructed person.

Some have expressed concern over the potential that big data has to make discrimination easier. For example, bigots could hide behind impersonal algorithms prebaked to exclude women or minorities. Others have expressed concern that even if not consciously used to discriminate, big data driven algorithms might nonetheless classify along racial or gender lines.³⁶ 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 offers, insurance rates, or employment opportunities.³⁷ 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 Article.

³⁵ See, e.g., Ohm, *supra* note 26, at 34 (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).

³⁶ See Crawford & Shultz, *supra* note 7, at 101; Elizabeth Dwoskin, *How Social Bias Creeps into Web Technology*, WALL ST. J., Aug. 21, 2015.

³⁷ One notable exception is the work of Lior Strahilevitz. He 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 disproportionately high imprisonment rates. See Lior Jacob Strahilevitz, *Privacy vs. Antidiscrimination*, 75 U. CHI. L. REV. 363, 376 (2008).

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.³⁸ 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.³⁹

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.”⁴⁰

Other privacy 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.”⁴¹ Target has become the poster child of classification harms for using analytics to send coupons for maternity-related products, such as diapers, cribs, and prenatal vitamins to customers whose shopping habits suggested a

³⁸ See Scott R. Peppet, *Regulating the Internet of Things: First Steps Towards Managing Discrimination, Privacy, Security, and Consent*, 93 TEX. L. REV. 85 (2014).

³⁹ *Id.* at 118.

⁴⁰ *Id.* at 126.

⁴¹ 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) (stating that, 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. ONLINE 35, 36 (2013) (noting “concerns with the classifications and segmentation produced by big data analysis”); Ira Rubinstein, *Big Data: The End of Privacy or a New Beginning*, 3 INT’L DATA PRIVACY L. 74 (2013).

high likelihood of being pregnant.⁴² 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 discounts.⁴³ 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.⁴⁴ 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.⁴⁵ 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.”⁴⁶ 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.⁴⁷

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

⁴² See, e.g., Crawford & Shultz, *supra* note 41, at 98–99; Harford, *supra* note 28.

⁴³ Despite widespread outrage, it is unclear whether the critics would prefer Target to provide prenatal vitamin discounts to everybody or to nobody. Peppet, *supra* note 7, at 1195.

⁴⁴ See Joseph W. Jerome, *Buying and Selling Privacy: Big Data's Different Burdens and Benefits*, 66 STAN. L. REV. ONLINE 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 Polonetsky, *Judged by the Tin Man: Individual Rights in the Age of Big Data*, J. 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 Polonetsky, *Big Data for All: Privacy and User Control in the Age of Analytics*, 11 NW. J. TECH. & INTELL. PROP. 239 (2013). In earlier work, Peppet explains how consumer signaling may substitute for firm screening as consumers, employees, and 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 will naturally infer that non-reporters are of the “bad” type. Peppet’s concern is that the increasing ability to credibly reveal one’s type will reduce privacy by raising the price of non-revelation. See Peppet, *supra* note 7. 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 41, at 37.

⁴⁵ Calo, *supra* note 32.

⁴⁶ *Id.* at 1133.

⁴⁷ *Id.* at 1131.

offer discounts to the rich on luxury goods, which are subsidized by high prices for the poor on staples such as bread and milk.⁴⁸ Crawford and Shultz lament the fact that rich and poor receive different credit offers online.⁴⁹ Further, Joseph Jerome concedes that big data will enhance market efficiency but nonetheless warns “market efficiency favors the wealthy, established classes.”⁵⁰ He adds that “categorization and classification threaten to place a privacy squeeze on the middle class as well as the poor.”⁵¹

Not surprisingly, authors generally recommend a government response to the problems posed by big data.⁵² 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.⁵³ 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, and “mixed uses” involving both advertising and health information, like the Target pregnancy debacle, would receive the same protection as health information.⁵⁴ 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.⁵⁵

⁴⁸ See Omer Tene, *Privacy: For the Rich or for the Poor*, CONCURRING OPINIONS (July 26, 2012), <http://concurringopinions.com/archives/2012/07/privacy-for-the-rich-or-for-the-poor.html>.

⁴⁹ See Crawford & Shultz, *supra* note 7, at 101.

⁵⁰ Jerome, *supra* note 44, at 50.

⁵¹ *Id.*

⁵² See, e.g., Peppet, *supra* note 38, at 150–56 (arguing for restrictions on “cross-context” use of data streams and analogizing them to FCRA, the 5th Amendment, and the Genetic Information Nondiscrimination Act); Dwork & Mulligan, *supra* note 41, at 39 (suggesting the establishment of a metric “defining who must be treated similarly” that “creat[es] a path for external stakeholders ... to have greater influence over, and comfort with, the fairness of classifications.”).

⁵³ Peppet, *supra* note 7, at 1158–59.

⁵⁴ Crawford & Shultz, *supra* note 41, at 118.

⁵⁵ *Id.* at 126–28.

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 calls "data determinism," which occurs when individuals are judged "because 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."⁵⁶ Her colleague, Commissioner Julie Brill, similarly has expressed concern that "[t]he 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."⁵⁷

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,⁵⁸ "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."⁵⁹ Based on these potential classifications, the FTC recommended Fair Credit Reporting Act (FCRA)-like legislation to cover data brokers.⁶⁰ Most recently, the FTC issued its Big Data Report, which addressed the potential for big data to act as a "tool for exclusion" of certain disadvantaged populations.⁶¹ The Report cautioned against the use of big data resulting in "individuals being ... denied opportunities" and "higher-priced goods and services for lower income communities," as

⁵⁶ See Edith Ramirez, Chairwoman, Fed. Trade Comm'n, *The Privacy Challenges of Big Data: A View from the Lifeguard's Chair*, Keynote Address at the Technology Policy Institute Aspen Forum (Aug. 19, 2013). See also Ramirez, *supra* note 7 (warning of using big data to segment along income or racial lines, and referring to this practice as "discrimination by algorithm" and "digital redlining").

⁵⁷ Brill, *supra* note 7.

⁵⁸ FED. TRADE COMM'N, *DATA BROKERS: A CALL FOR TRANSPARENCY AND ACCOUNTABILITY* 48 (2014).

⁵⁹ *Id.*

⁶⁰ *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.

⁶¹ FED. TRADE COMM'N, *BIG DATA: A TOOL FOR INCLUSION OR EXCLUSION* (2015), available at <https://www.ftc.gov/system/files/documents/reports/big-data-tool-inclusion-or-exclusion-understanding-issues/160106big-data-rpt.pdf>.

well as the possibility of “[c]reat[ing] or reforc[ing] existing disparities,” and “expos[ing] “sensitive information.””⁶² Although the report did not recommend legislation, it laid out best practices for firms that use big data, including making sure not to confuse correlation with causation and taking steps to eliminate “hidden biases” and “unintended impacts on certain populations.”⁶³

Finally, the White House recently floated a draft privacy bill that adopted a strong regulatory stance toward big data predictions.⁶⁴ 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.⁶⁵

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.⁶⁶ 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

At the end of the day, those concerned with classification harms really are concerned with big data’s potential to promote separating equilibria.⁶⁷ Such concerns, however, run contrary to the general proposition that separation is better than pooling. Because it

⁶² *Id.* at 9–11.

⁶³ *Id.* at 32.

⁶⁴ Office of Management & Budget, ADMINISTRATION DISCUSSION DRAFT: CONSUMER PRIVACY BILL OF RIGHTS ACT OF 2015 (2015).

⁶⁵ *See id.* at § 103.

⁶⁶ *See* Peppet, *supra* note 7; Jerome, *supra* note 44, at 51.

⁶⁷ *See* Peppet, *supra* note 7, at 1188–90.

lies at the heart of the matter, it is useful to explore these concepts in some detail.

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*.⁶⁸ There are two types of millers—those with a spare shaft (good types) who will continue to operate when one breaks and those without (risky 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. Those with a spare shaft 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 spare-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.⁶⁹ This price is a bargain for one-shafters; they 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

⁶⁸ *Hadley v. Baxendale*, 156 Eng. Rep. 145 (Ex. Ch. 1854).

⁶⁹ Ian Ayers & Robert Gertner, *Filling Gaps in Incomplete Contracts: An Economic Theory of Default Rules*, 99 *YALE L.J.* 87 (1989).

adverse selection. Risky types are drawn into the market because they can free-ride off of good types. Of course, this causes good types to leave, resulting in a market characterized by lower output and a greater proportion of risky types than would exist with full information. In the extreme, adverse selection can cause markets to unravel completely.⁷⁰

Adverse selection can be found in a variety of markets in which one party is likely to have private information.⁷¹ 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.⁷² Two potential employees may look similar

⁷⁰ George Akerlof, *The Market for Lemons: Quality Uncertainty and the Market Mechanism*, 84 Q. J. ECON. 488 (1970).

⁷¹ 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. MONETARY ECON. (2006) (consumer loan market); Liran Einav, Mark Jenkins, & Jonathan Levin, *The Impact of Credit Scoring on Consumer Lending*, 44 RAND J. ECON. 249 (2013) (subprime auto loan market); William Adams, Liran Einav, & Jonathan Levin, *Liquidity Constraints and Imperfect Information in Subprime Lending*, 99 AM. ECON. REV. 49 (2009) (subprime auto loan market); Bev Dahlby, *Testing for Asymmetric Information in Canadian Automobile Insurance*, in CONTRIBUTIONS TO INSURANCE ECONOMICS 423 (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 ‘Unused Observables’ in Insurance Markets: Evidence from the U.K. Annuity Market*, 81 J. RISK & INS. 709 (2014); Dean Karlan & Jonathan Zinman, *Expanding Credit Access: Using Randomized Supply Decisions to Estimate the Impacts*, 23 REV. FIN. STUD. 433 (2010) (South African subprime lender); Robery Puelz & Arthur Snow, *Evidence on Adverse Selection: Equilibrium Signaling and Cross-Subsidization in the Insurance Market*, 102 J. POL. ECON. 236 (1994). But see Pierre-Andre Chiappori & Bernard Salanie, *Testing for Asymmetric Information in Insurance Markets*, 108 J. POL. ECON. 56 (2000) (finding 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).

⁷² This difference is the point behind esurance’s “sorta you isn’t you” campaign. See ESURANCE, https://www.esurance.com/quote1301?PromoID=GGNBB_VA_001&partner_cd=AdPos-1t1%7CGeo-9008162%7CAdID-79094026107%7C&ts=2 (last visited Apr. 14, 2017).

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. Further, 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.⁷³ In credit markets, lenders understand that higher interest rates will attract a disproportionate share of consumers who are more likely to default. Employers who offer lower wages risk attracting only the least productive workers.

Firms ideally would like to find a way to separate good from risky types and offer each a contract that reflects their true types. One strategy is to screen potential customers by offering a set of contracts that will create incentive for types to reveal themselves. This strategy 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.⁷⁴ If the contract offered to good types is too favorable, it will attract both types and prevent separation.⁷⁵ 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, and caps—to clear markets. As a result, good types bear too much risk or receive too little credit compared to the full information equilibrium.

⁷³ 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 consume more health care than they otherwise would. See Karlan & Zinman, *supra* note 71; Einav, Jenkins, & Levin, *supra* note 71; Finkelstein & Poterba, *supra* note 71.

⁷⁴ See Joseph E. Stiglitz & Andrew Weiss, *Credit Rationing with Imperfect Information*, 71 AM. ECON. REV. 393 (1981).

⁷⁵ This example is analogous to a test that has only easy questions, allowing the poor students to pool with the good.

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 cannot merely be declaring themselves good. This statement is “cheap talk”—a signal that is 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.⁷⁶

Clearly, in markets characterized by adverse selection, risky 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 risky 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, 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

⁷⁶ 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 & Marianna E. Page, *Degrees Matter: New Evidence on Sheepskin Effects in the Returns to Education*, 78 REV. 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. POL. ECON. 985 (1974).

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 can 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 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.⁷⁷

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 to incentivize 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.⁷⁸ Edleburg, moreover, finds evidence of moral hazard in U.S. consumer lending markets. Several papers have also found evidence of moral hazard in insurance markets.⁷⁹ Moral hazard

⁷⁷ See Chiappori & Salanie, *supra* note 71, at 60.

⁷⁸ Karlan & Zinman, *supra* note 71. See also Edleburg, *supra* note 71 (finding evidence of moral hazard in auto and credit card lending in the U.S.).

⁷⁹ See Puelz & Snow, *supra* note 71; Yingying Dong, *How Health Insurance Affects Health Care Demand—A Structural Analysis of Behavioral Moral Hazard and Adverse Selection*, 51 *ECON. INQUIRY* 1324 (2011). See also Einav, Jenkins & Levin, *supra* note 71; Jonathan Klick & Thomas Stratmann, *Subsidizing Addiction:*

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 product liability.⁸⁰ Finally, several studies document the so-called “Peltzman” effect, in which actors take less care when there are exogenous increases in safety.⁸¹

III. BIG DATA

It is not hard to see how big data could improve the performance of markets fraught with asymmetric information. Some may wonder just how much more companies can learn more about us: Amazon knows our purchase history; Netflix knows what we watch; and Google knows what we search for. Nonetheless, big data can allow more granular predictions, which, at least at the margin, will allow for better matching and amelioration of problems of adverse selection, as well as problems associated with moral hazard by more closely aligning outcomes with effort. Further, these benefits are likely to accrue more sharply to those at the lower rungs of the economic ladder. For example, because the poor tend to interact less frequently with the traditional credit system, the ability to combine disparate pieces of information—from phone bills to social media

Do State Health Insurance Mandates Increase Alcohol Consumption?, 35 J. LEGAL STUD. 175 (2006) (finding evidence of moral hazard when state laws mandate coverage of diabetes and alcohol abuse treatment). *But see* Chiappori & Salanie, *supra* note 71 (finding no evidence of adverse selection in French automobile insurance market).

⁸⁰ Paul H. Rubin & Joanna M. Shepherd, *Tort Reform and Accidental Deaths*, 50 J. L. & ECON. 221 (2007). Similarly, Helland & Taborock 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).

⁸¹ 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. POL'Y ANALYSIS & MGMT. 270 (1993).

posts—to form a more accurate picture will allow those with relatively better risk profiles to reveal themselves as such.

A. Adverse Selection

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. By allowing finer segmentation of risk pools, it will reduce pooling equilibria, which reduce welfare by causing good types to exit markets. Further, as big data makes screens more accurate, it will reduce the need for costly signals, such as collateral or education, to effect separation.

For example, alternative credit scoring mechanisms use a variety of predictors, such as social media posts and payment of cell phone bills, to predict credit worthiness.⁸² Individuals with stable networks of close friends and whose information on LinkedIn matches his application or businesses with good reputations on social media are more likely to get loans.⁸³

Similarly, some employers are using big data predictions about potential employees to supplement, or even replace, traditional hiring techniques for some jobs.⁸⁴ Some companies are beginning to use big data analytics to identify candidate employees for technology, high-end sales, and managerial positions, and these analytics are suggesting that other indicators are more predictive of good fits than college.⁸⁵ In

⁸² Nate Cullerton, *Behavioral Credit Scoring*, 101 GEO. L.J. 808, 809 (2013); 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*, POL. & ECON. RES. COUNCIL (2012); Elizabeth Dwoskin, 'Big Data' Doesn't Yield Better Loans, WALL ST. J., Mar. 17, 2014; 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; *Crunching the Numbers*, ECONOMIST, May 19, 2012, at 7.

⁸³ See Stephanie Armour, *Borrowers Hit Social-Media Hurdles*, WALL ST. J., Jan. 8, 2014.

⁸⁴ Claire Cain Miller, *Can an Algorithm Hire Better than a Human?*, N. Y. TIMES, June 25, 2015; Max Nisen, *MONEYBALL AT WORK: They've Discovered What Really Makes a Great Employee*, BUS. INSIDER, May 6, 2013.

⁸⁵ See Don Peck, *They're Watching You at Work*, THE ATLANTIC Dec. 2013; See also supra notes 82–86 and accompanying text.

addition to facilitating better matching between wages and workers, this use of big data also has the potential to reduce socially wasteful signaling expenditures. As discussed in Part II.A, a large body of empirical work lends support to the signaling value of education.⁸⁶ To the extent that these new analytic techniques can save individuals from making larger investments in education than they otherwise would, they are socially beneficial. Further, one firm has examined employee email, calendars, and HR records, and found a correlation between attendance of events and benefits coverage selection and the likelihood of an employee quitting within a year.⁸⁷ 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.⁸⁸

B. Moral Hazard

What role might big data play in ameliorating problems associated with moral hazard? Moral hazard can be tempered through separation. Risky types lose in separating equilibria because good types no longer subsidize them. However, in some cases, types are not immutable characteristics but instead are a result of choices made under moral hazard. Recall the separation brought about by the court in *Hadley v. Baxendale*.⁸⁹ Although the one-shaft millers are worse off, society gains because pricing sends more accurate signals that tell the courier to take more care with one-shaft millers. These higher prices also mean that the one-shaft miller now bears 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, 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 risky types forces them to pay prices that more closely reflect their true risk, those at the margins will alter their behavior to the extent that it is feasible. For example, if alternative credit scoring limits pooling among subprime

⁸⁶ See *supra* note 76 and accompanying text.

⁸⁷ Rachel Emma Silverman & Nikki Waller, *The Algorithm that Tells the Boss Who Might Quit*, WALL ST. J. Mar. 13, 2015.

⁸⁸ *Id.*

⁸⁹ *Hadley v. Baxendale*, 156 Eng. Rep. 145 (Ex. Ch. 1854).

populations, it may incentivize those with 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.

C. Benefits for the Poor

At this point, it is worth a minor detour to explain why, contrary to the worries of several privacy law scholars,⁹⁰ there are strong reasons to believe that benefits from big data-driven separation may accrue disproportionately to the poor. It does not take big data to tease out large differences—say between prime and subprime markets. However, within a population, the more granular sorting made possible by big data can help separate out those good types who have been pooled with their relatively riskier cohorts. In this manner, big data analytics are likely to be especially valuable to those consumers for whom there is little traditional information that can be used for screening. Further, as big data allows firms to achieve more accurate segmentation through screening, it can reduce the need to rely on signals.⁹¹ Again, this benefit is likely to accrue to those who cannot afford to send costly signals.

First consider the empirical evidence on the impact of credit scoring and risk-based pricing—in many ways, the precursor to today's big data—on the poor's access to credit, which is suggestive of big data's potential to have a positive impact on those on the lower rungs of the economic ladder. For example, Federal Reserve data show that from 1983 to 2010, the largest increases in credit card

⁹⁰ See, e.g., Tene, *supra* note 48 (discounts to the rich subsidized by price hikes for the poor); Jennifer Valentino-DeVries, Jeremy Singer-Vine, & Ashkan Soltani, *Websites Vary Prices, Deals Based on Users' Information*, WALL ST. J., Dec. 24, 2012) (finding that differential online pricing based on zip code leads to those in relatively poorer zip codes to pay more; FED. TRADE COMM'N, *supra* note 90 at 48 (expressing concern that the poor will be marketed only subprime offers); Ramirez, *supra* note 7, at 4 (expressing concern over "data determinism," which will limit options for the poor).

⁹¹ See *infra* notes 100–06 and accompany text.

ownership are in the bottom half of income earners (200–300%).⁹² Moreover, from 1970 to 2010, there was a 77% increase in access to consumer credit by the lowest quintile compared to a 14% increase for the highest quintile.⁹³ The poor also appeared to gain from automated underwriting for mortgages. For example, a study finds that automated underwriting (AU) based on credit scoring is more accurate at predicting risk than manual underwriting and as a result approves more lower-income borrowers.⁹⁴ The authors conclude:

It is not surprising that the increased accuracy of AU 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.⁹⁵

A 2006 Federal Reserve Board report to Congress on credit scoring echoed these themes.⁹⁶ 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.”⁹⁷ 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 to 2004 and that in any event, there was

⁹² See THOMAS A. DURKIN, GREGORY ELLIEHAUSEN, MICHAEL E. STATEN, & TODD J. ZYWICKI, CONSUMER CREDIT AND THE AMERICAN ECONOMY 304 (2014).

⁹³ *Id.* at 72.

⁹⁴ See Susan Wharton Gates, Vanessa Gail Perry, & Peter M. Zorn, *Automated Underwriting: Good News for the Underserved?* 13 HOUSING POL'Y DEBATE 369 (2002).

⁹⁵ *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 \$0.5 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).

⁹⁶ 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 (2007).

⁹⁷ *Id.*

no evidence that those in the top of the income distribution disproportionately gained from increased information about credit history.⁹⁸

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.⁹⁹ A used car seller dealt with a very financially distressed population: average annual income was \$28,000 and default rates on loans were over 60%. Further, there was strong evidence of both moral hazard and adverse selection: 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 riskiest borrowers about twenty percentage points more likely to default than the least risky.¹⁰⁰ According to the authors, these data allowed the dealer to separate “consumers with transitory bad records from persistently bad risks.”¹⁰¹ 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.¹⁰² Many lower-income consumers have little or no information on file with

⁹⁸ *Id.*

⁹⁹ Einav, Jenkins, & Levin, *supra* note 71.

¹⁰⁰ *Id.*

¹⁰¹ *Id.* at 255.

¹⁰² 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*, POL. & ECON. RES. COUNCIL (2012).

credit reporting agencies. As a result, lenders are unable to make reliable inferences about them: “unscorable” consumers are 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 than they typically deserve.¹⁰³ 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%, compared to only a 5% increase for the highest income group.¹⁰⁴ 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.¹⁰⁵

Today, several start-ups are using big data to analyze thousands of variables like rent records, prior payday loans, pawnshop transactions, and Facebook friends to identify better credit risks within poor populations.¹⁰⁶ 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.¹⁰⁷ The upshot is that these alternative scoring systems can give underserved populations alternatives to payday lenders or pawnshops, and by identifying creditworthy

¹⁰³ *See id.* at 13.

¹⁰⁴ *Id.* at 17 (Figure 7).

¹⁰⁵ *Id.*

¹⁰⁶ *See* Elizabeth Dwoskin, ‘Big Data’ Doesn’t Yield Better Loans, WALL ST. J, Mar. 17 2014.

¹⁰⁷ *See* Jenkins, *supra* note 82; *see also*, Quentin Hardy, *Big Data For the Poor*, N.Y. TIMES, July 5, 2012 (suggesting that although these alternative credit outlets offer triple-digit interest rates due to high default rates, they still claim to offer cheaper alternatives than payday loans for subprime populations).

individuals within a population of high credit risks, lenders are reducing default rates below those experienced by payday lenders.¹⁰⁸

Benefits of using big data in employment also are likely to accrue in large portion to those on the lower rungs of the economic ladder—those who can least afford to send signals by purchasing postsecondary education. The data are clear that the largest share of economic gains over the past three decades have gone to those with college degrees.¹⁰⁹ At the same time, many entry-level jobs require postsecondary 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. As noted above, some companies are beginning to use big data analytics to find potential employees for positions that typically require a college degree or at least some postsecondary education.¹¹⁰ To the extent that big data can unlock the doors to jobs that were previously only for those with a college education, it may allow a wider sharing of economic gains.¹¹¹

Finally, consider concern voiced by several authors that big data-driven price discrimination will cause the poor to pay more than the rich.¹¹² 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.¹¹³ 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

¹⁰⁸ Big Data: Crunching the Numbers, THE ECONOMIST, May 19, 2012.

¹⁰⁹ See Bureau of Labor Statistics, *Labor Force Statistics from the Current Population Survey*, UNITED STATES DEPARTMENT OF LABOR, BUREAU OF LABOR STATISTICS (Apr. 20, 2017) <https://www.bls.gov/cps/earnings.htm#education>.

¹¹⁰ See Don Peck, *They're Watching You at Work*, THE ATLANTIC (Dec. 2013), available at <https://www.theatlantic.com/magazine/archive/2013/12/theyre-watching-you-at-work/354681/>.

See also *supra* notes 82–86 and accompanying text.

¹¹¹ See Josh Zumbrun, *Just How Stagnant are Wages Anyway?*, WALL ST. J., Jul. 6, 2015.

¹¹² Calo, *supra* note 22, at 1029 (stating 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”).

¹¹³ See, e.g., Jerome, *supra* note 44 (suggesting we know that poor rarely win from competition); see also FED. TRADE COMM’N, *supra* note 61.

becomes easier to implement.¹¹⁴ Assertions that price discrimination brought about by big data is likely to allow firms to implement schemes under which the poor 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.¹¹⁵ 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,¹¹⁶ for example, consider student or elderly discounts at movies and restaurants or the Saturday stay-over and advance booking requirements for cheaper flights.¹¹⁷ 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 Mac users under the assumption that Mac users typically are wealthier than PC users.¹¹⁸

Adverse selection and moral hazard are real problems that force good types to subsidize risky types and reduce resources available to society as whole. Big data has the potential to ameliorate these problems by allowing accurate screening mechanisms to reveal more granular distinctions within distributions. Moreover, these distinctions have great potential to benefit disadvantaged populations by identifying relatively good risks within a larger cohort of risky types. However, ameliorating information asymmetries often will

¹¹⁴ See EXECUTIVE OFFICE OF THE PRESIDENT, DIFFERENTIAL PRICING 17 (2015).

¹¹⁵ This calculation is called Ramsey pricing, and formally requires:

$PAPB = 1 + 1/\epsilon_A + 1/\epsilon_B$, 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).

¹¹⁶ 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. DRUG ISSUES 631, 635 (1998).

¹¹⁷ See N. GREGORY MANKIW, PRINCIPLES OF MICROECONOMICS (6th ed. 2011).

¹¹⁸ 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. Further, as a recent White House report noted, there is little evidence to suggest that firms are engaging in the practice. See EXECUTIVE OFFICE OF THE PRESIDENT, DIFFERENTIAL PRICING 10 (2015).

implicate privacy. In the next Part, this Article develops a framework that helps distinguish productive from dissipative uses of big data.

IV. SEPARATING BENEFICIAL FROM HARMFUL USES OF BIG DATA

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 is 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. Additionally, information collection itself can be dissipative.

In this section, I draw on the economic theory of contracts and torts to develop a framework that identifies circumstances in which privacy is likely to be either socially wasteful or beneficial. I show that any sensible policy toward big data should attempt to avoid promoting privacy or forced disclosure that serves *only* to make one party better off at the expense of the other. In some—perhaps most—instances, using big data to make predictions will come with both social benefits and privacy costs. I provide a mechanism for balancing separation benefits against intrinsic privacy harms 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) about themselves. Consumers with higher R s are better types (*e.g.*, better credit or insurance risks or better employees). Good types would like to reveal themselves but cannot do so credibly. Firms (collectively referring to insurers, employers, lenders, and producers), however, can use big data to screen good types from risky types by revealing 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 surplus (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: $RV(x) - x$.¹¹⁹ 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.¹²⁰ For example, those with higher R s will receive better credit or employment offers than risky 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} .¹²¹ 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.¹²² Because \bar{x} is profit-maximizing only for \bar{R} , $V(x)$

¹¹⁹ This framework is similar to that used in, Steven Shavell, *Acquisition and Disclosure of Information Prior to Sale*, 25 RAND J. ECON. 20 (1994).

¹²⁰ The first order condition for profit maximization requires that $RV_x(x^*(R)) - 1 = 0$. Differentiating this expression with respect to R yields:

$$\frac{\partial x^*}{\partial R} = \frac{-V_x}{RV_{xx}} > 0.$$

¹²¹ 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 R for all types.

¹²² I have normalized the action for the lowest value of R to zero so that $x^*(R)$ is positive. For example, if R measures productivity or credit worthiness, x would measure wages or credit terms, with higher values of x representing higher wages or more attractive credit terms, respectively. In Figure 1, $x^*(R)$ is linear to reflect the assumption that marginal benefit from matching does not vary with R . There could

is lower than it could be when firms must take \bar{x} for any type not equal to \bar{R} . Accordingly, the gap between $x^*(R)$ and $\bar{x}(\bar{R})$ represents the social loss from pooling versus separating equilibria. Further, these social losses may be exacerbated by moral hazard. In a separating equilibrium, consumers will have incentives to take actions to improve their type to gain a larger payout. For example, use of big data to predict diabetes risk or driving ability may spur healthier lifestyles or safer driving habits. Not only are taking such actions privately rational, they improve social welfare by increasing the average population type.

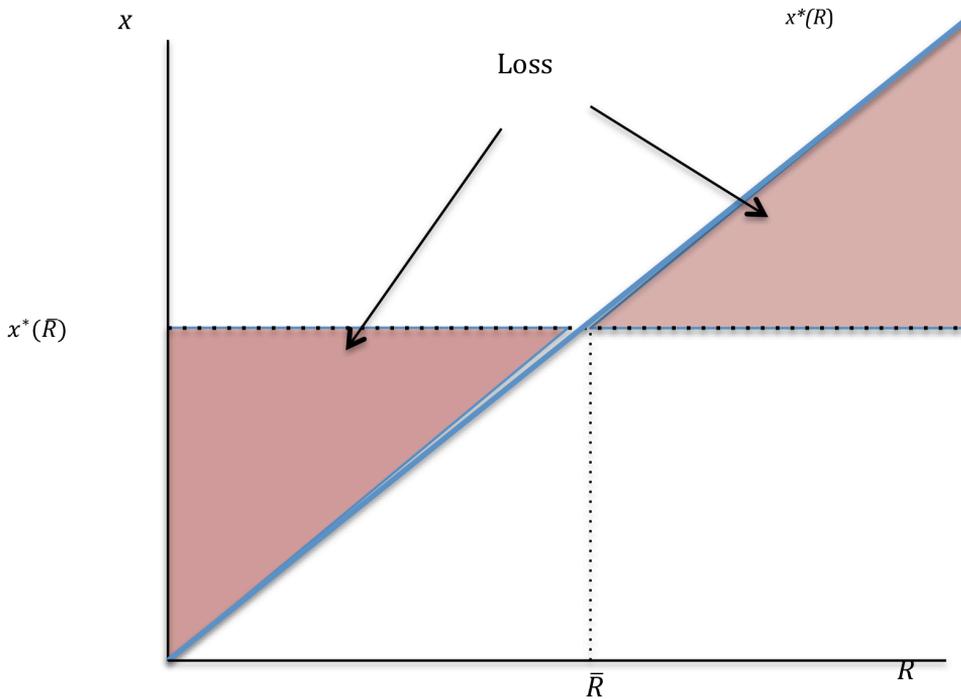


FIGURE 1. OPTIMAL ACTION BY TYPE

Consumers receive two potential payoffs. If big data is used to match action with types, they receive:

be situations in which the marginal benefits vary at extreme levels, in which case $x^*(R)$ would be non-linear. But the inclusion of such instances would not impact the analysis.

$$U = x_i^*(R_i) - t,$$

where $x_i^*(R_i)$ is the action for consumer i 's specific type. Utility increases in x , which represents more favored actions due to higher R s (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. In a big data regime, consumers also lose t , which is a transfer to firms with big data-driven knowledge of their reservation prices. If big data cannot be used for separation, consumers receive:

$$U = \bar{x}(\bar{R}) + P,$$

where $\bar{x}(\bar{R})$ is the uniform action based on the average type.

In this regulatory regime, consumers also receive a privacy benefit, P , from the ability to prevent big data from classifying them.¹²³ This intrinsic value of privacy captures several dimensions of privacy value unrelated to bargaining gains from concealment.¹²⁴ For example, people clearly value being free from unwanted observation and intrusions into their private spheres, although this value varies across the population and contexts.¹²⁵ For example, P can capture the value of not being tracked online (the input into a big data algorithm) or not being sent targeted ads (the output of a big data algorithm). 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 dull incentives to become tested in the first place,

¹²³ The focus of intrinsic harm is on classification that reveals something the consumer would prefer to keep private. It also could involve the collection of sensitive data used to feed the algorithm. In the context of big data, however, the privacy concerns often relate to the use of myriad bits of non-sensitive data to assemble a picture of a sensitive trait. See FED. TRADE COMM'N, *supra* note 61, at 10.

¹²⁴ See Richard S. Murphy, *Property Rights in Personal Information: An Economic Defense of Privacy*, 84 GEO. L.J. 2381, 2385 (1996) (discussing "privacy divorced from reputation," as distinguished from privacy, which is merely fraud to enhance one's reputation).

¹²⁵ See *infra* notes 153–63 and accompanying text.

although such knowledge clearly is valuable.¹²⁶ Just as copyrights and patents are designed to foster incentives to create and invent, moreover, providing exclusive rights in personal information can enhance incentives for self-discovery.¹²⁷ Finally, ubiquitous surveillance and predictions from the resulting data can lead to wasteful 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.

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 to see in Table 1 how risky types lose in a separating equilibrium.¹²⁸ The impact of big data revelation is unambiguously negative for risky types ($R_R < \bar{R}$). They gain from revelation only if $(x_R^* - \bar{x}) > P + t$, which can never hold because $(x_R^* - \bar{x}) < 0$. Good types ($R_G > \bar{R}$) will prefer privacy to revelation only if their intrinsic value of privacy is greater than their share of increased in surplus from reductions in adverse selection (net of any pure transfers to firms): $[(x_G^* - \bar{x}) - t < P]$.¹²⁹ Thus, demand to restrict big data predictions is driven by a coalition of risky types who stand to gain from strategic concealment and good types who are privacy sensitive.

¹²⁶ 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).

¹²⁷ See Murphy, *supra* note 124, at 2386-87. To the extent that there are positive externalities for society that are not captured in P —for example those that underlie rights-based notions of privacy, which focuses on notions of autonomy that are necessary to spur the type of diversity, creativity, and intellectual development that serves society as a whole—private incentives may not be optimal. 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); Julie Cohen, *What Privacy Is For*, 126 HARV. L. REV. 1904, 1911 (2013); Neil Richards, *Intellectual Privacy*, 87 TEX. L. REV. 387, 407 (2008).

¹²⁸ I assume that signaling is too expensive for good types in this simple model, so that separation cannot be achieved through self-selection.

¹²⁹ 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): $(\partial U/\partial x)/(\partial U/\partial P)$.

TABLE 1. CONSUMER PAYOFFS IN PRIVACY AND REVELATION REGIMES

		<i>Legal Regime</i>	
		Privacy	Revelation
<i>Consumer Type</i>	Good	$\bar{x}(R_G) + P$	$x_G^*(R_G) - t$
	Risky	$\bar{x}(R_R) + P$	$x_R^*(R_R) - t$

B. Dissipative Privacy and Dissipative Revelation

As shown above, privacy is intrinsically valuable, but it's also privately valuable when concealment of relevant information leads to better terms of trade—what I refer to as “strategic concealment.” A key difference between these types of privacy is that one creates value and the other destroys it. Strategic privacy is purely dissipative; although it's privately rational to want to conceal information that will reduce your share of surplus from a bargain, such privacy is socially wasteful because it deprives society of the gains from reducing asymmetric information. Thus, 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 known that there can be socially excessive incentives to collect information; over forty years ago, Hirshleifer 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.¹³⁰ Otherwise, knowledge serves only to redistribute surplus, and hence expenditures to collect it are dissipative. Any sensible

¹³⁰ Jack Hirshleifer, *The Private and Social Value of Information and the Reward to Inventive Activity*, 61 AM. ECON. REV. 561 (1971).

policy toward big data, therefore, should attempt to avoid promoting privacy or disclosure that serves *only* to move surplus from one party to another.

A paradigm for this framework can be found in contract law. For example, sellers typically are required to disclose unfavorable information about their wares.¹³¹ 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.¹³² 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.¹³³ 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 foreknowledge of conditions that impact the value of a commodity.¹³⁴ The distinction between duress and necessity also has an economic rationale 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.¹³⁵ 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.¹³⁶ Finally, the limitation on consequential damages creates incentives for buyers to reveal private information about their sensitivity to breach.¹³⁷ Concealment in these

¹³¹ 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).

¹³² See Larry E. Ribstein & Bruce H. Kobayashi, *State Regulation of Electronic Commerce*, 51 *EMORY L.J.* 1, 15 (2002).

¹³³ 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).

¹³⁴ See *Laidlaw v. Oregon*, 15 U.S. 178, 194–95 (1817); STEVEN SHAVELL, *FOUNDATIONS OF ECONOMIC ANALYSIS OF LAW*, 332–35 (2006).

¹³⁵ See ROBERT COOTER & THOMAS ULEN, *INTRODUCTION TO LAW AND ECONOMICS*, 281–87 (6th ed. 2011).

¹³⁶ *Id.*

¹³⁷ See *Hadley v. Baxendale*, 156 Eng. Rep. 145 (Ex. Ch. 1854); Ayers & Gertner, *supra* note 69, at 104–05.

circumstances is 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 Privacy

For privacy to be dissipative, three conditions must be met. First, concealment of the information must retard value-creating actions—actions that reduce losses from adverse selection. Second, the value that would be created from these actions must be greater than the cost of discovering the private information. That is, the use of big data must pass a benefit-cost test. Finally, the only gains from concealment must be strategic.¹³⁸ If these conditions are met, privacy is dissipative because the only gains from concealment come in the form of an increased share of surplus to risky types at the expense of total surplus. This case is easily shown. For example, consider a world with two workers—one productive and one lazy. If there are no intrinsic privacy gains from concealing one's laziness, privacy is welfare enhancing only if:

$$(\bar{x} - x_R^*) > (x_G^* - \bar{x}) + \Delta V.$$

The left-hand side of this expression is the gain to lazy types from concealment—the subsidy they receive from productive types. The first part of the right-hand side is the gain to productive types from revelation—the amount they subsidize lazy types in a pooling equilibrium. Finally, the term ΔV is the increase in surplus due to better matching.¹³⁹ So a necessary (but not sufficient) condition for purely strategic privacy to be socially beneficial is that gains to risky types must outweigh gains to good types as we move from a separating to a pooling equilibrium. This condition implies that the following inequality must hold:

¹³⁸ If concealment gives rise to intrinsic privacy gains that are less than the value created from reducing asymmetric information, privacy is not dissipative—because it creates value—just socially inefficient.

¹³⁹ $\Delta V = [(R_G(V(x_G^*) - V(\bar{x})) - x_G^* + \bar{x}) + (R_B(V(x_R^*) - V(\bar{x})) - x_R^* + \bar{x})] > 0$. This is due to the fact that $RV(x^*) - x^* \geq RV(\bar{x}) - \bar{x}$ for all types, because x^* maximizes $R(V(x) - x)$.

$$\bar{x} > \frac{x_G^* + x_R^*}{2}.$$

But this condition can never hold because $\bar{x} = \frac{x_G^* + x_R^*}{2}$. Thus, as long as $\Delta V > 0$, strategic concealment will always reduce welfare because gains to risky types are bounded from above by losses to good types.

To make this result more concrete, consider the potential employee whose expected productivity score from an accurate big data algorithm is too low to garner an interview.¹⁴⁰ The prediction that he's not a correct match for employment at this firm is disappointing to the applicant because he is no longer able to disguise his poor work habits by blending in with more productive workers. Of course, the only harm from unmasking his true type is the surplus that he would have extracted from the firm (from paying too high a wage) and the productive workers (from subsidizing his laziness by receiving too low a wage). Further, because the ability to sort productive from unproductive workers raises productivity, losses to this unproductive worker from separation is less than the gains to society as a whole, which include increased firm profits and increased utility to more productive workers whose market opportunities previously were limited due to adverse selection.¹⁴¹

Put differently, when privacy serves purely strategic purposes, losses to risky 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 more efficient matching of action with type. Without these losses, the net gains to society cannot materialize.¹⁴² Here, private information about the worker is put to its most valuable use when it's revealed.

¹⁴⁰ See, e.g., *supra* notes 82 and 83.

¹⁴¹ These gains may also include reductions in moral hazard from choosing to engage in y .

¹⁴² 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).

2. Dissipative Privacy and Antidiscrimination

It must be conceded at this point that there are situations in which using big data to ferret out useful information would increase surplus and harm no privacy interests, but forced pooling could be efficient for reasons outside the model. For example, consider a restaurant in a rural town that has a large population who is opposed to homosexuality on religious grounds. It may well be that using big data to predict sexual preference of its employees correctly will increase productivity;¹⁴³ people in this town may be less likely to frequent a restaurant with homosexual staff, or the restaurant may have difficulty hiring staff who want to work with homosexuals.

This hypothetical brings the issue of strategic privacy into sharp relief. Concealing sexual orientation (or race, religion, and gender) to avoid suffering bad marketplace outcomes is every bit as strategic as concealing one's poor history of repayment or gainful employment.¹⁴⁴ Yet, we have a legal framework in place to avoid classification based on the former set of traits and not the latter. When we recognize that allowing strategic privacy—and hence condoning pooling—is about discrimination and not privacy, it helps clarify the domains in which we're comfortable discriminating. Most modern societies would not tolerate firms offering different wages or credit terms based on race, religion, gender, and increasingly sexual orientation. At the same time, it would be hard to conceive of an antidiscrimination policy that forbids insurers from discriminating based on driving records or creditors discriminating based on credit history. Thus, defining the domain of traits on which discrimination will not be tolerated necessarily maps out its complement: the set of traits on which discrimination is allowable and hence the domain of dissipative privacy.

In addition to forcing society honestly to confront the implications of countenancing certain types of strategic privacy, relying on antidiscrimination law in these cases has two main advantages over a privacy regime. First, it discourages wasteful investments in signaling. If a firm is forbidden from making an

¹⁴³ See e.g., Michal Kosinski et al., *Private Traits and Attributes Are Predictable from Digital Records of Human Behavior*, 110 Proc. NAT'L ACAD. SCI. 5802, 5804 (2013).

¹⁴⁴ This assumes that there is no intrinsic privacy interest in concealment.

economic decision based on the trait in question, signaling no longer has value. Second, if revelation of the trait in question no longer leads to worse treatment, absent privacy concerns, consumers will invest less in concealment.

3. Dissipative Revelation

If big data predictions do not prompt surplus-enhancing actions, they are purely dissipative: firms are merely spending resources to transfer surplus from consumers to themselves.¹⁴⁵ This state is true whether or not consumers gain intrinsic value from concealment. When information collection costs more than the value it creates, privacy is always the most efficient policy as it preserves value.

Consider the following modification to the job-matching algorithm hypothetical from above: rather than predicting productivity, the employer uses big data to predict political leanings.¹⁴⁶ It's hard to conceive a circumstance in which this information may have bearing on productivity, so we can assume that big data is being used here for distributional purposes.¹⁴⁷ Thus, absent any output effects, infringing on privacy in this manner has no social value. A similar case can be made for the use of big data to create so-called "sucker lists" of vulnerable consumers who are likely to fall prey to scams. This investment creates no social value and, like expenditures on rent-seeking, serves only to transfer surplus from consumers to predatory firms at a cost to society. Here again, information is most valuable to society when it remains concealed.¹⁴⁸

Categorizing these types of predictions as privacy harms serves to preserve surplus in three ways: discouraging firms from expending resources to discover this type of information; discouraging consumers from spending resources to try to conceal this information; and eliminating any direct privacy harms from its revelation.

¹⁴⁵ Recall that if data collection provides no increase in value, it would be rational to collect information only if $t > 0$.

¹⁴⁶ See e.g., Michael Kosinski et al., *supra* note 143.

¹⁴⁷ For example, the employer may feel that it can get larger wage concessions from those with more liberal or socialist viewpoints.

¹⁴⁸ See David C. Vladeck, *Digital Marketing, Consumer Protection, and the First Amendment: A Brief Reply to Professor Calo*, 82 GEO. WASH. L. REV. ARGUENDO, 156, 162 (2014).

C. Concealment and Revelation Both Valuable

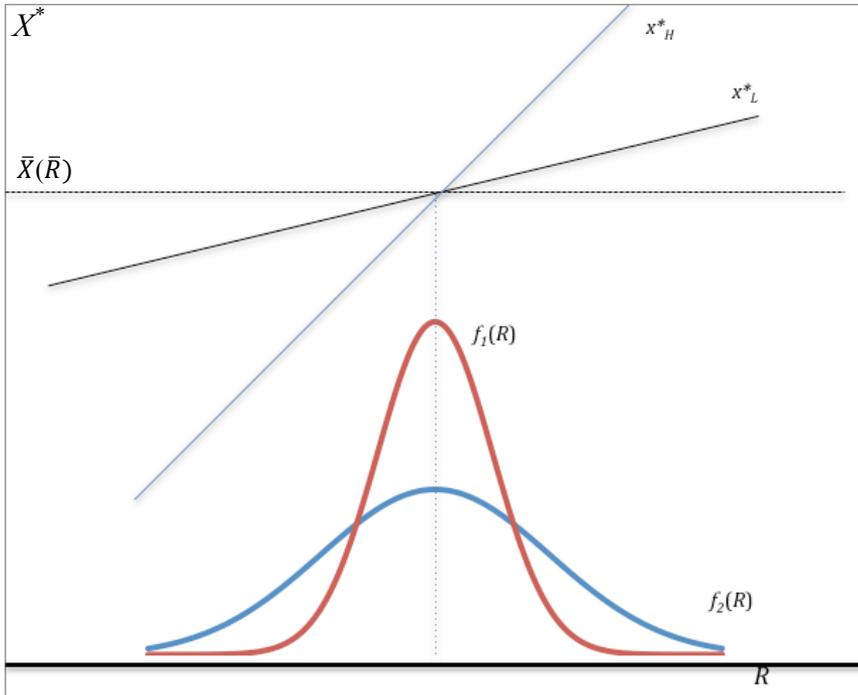
Things become more complicated when the discussion turns to sensitive data with both intrinsic and strategic value. In these cases, neither revelation nor privacy are dissipative: firms can increase surplus with big data predictions, but consumers also gain utility from concealment. In this section, 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 functions relating types to actions: x^*_L and x^*_H . Recall that in a privacy regime firms take a uniform action, $\bar{x}(\bar{R})$, for all types 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 each type ($x^*_i(R_i)$). It's easy to see that welfare losses from privacy are larger the steeper is x^* , reflecting the fact that the optimal action varies a great deal over types. For example, suppose x represents the premium a health insurer charges. We would expect x^* to be steeper if R measured propensity for substance abuse than if it measured shoe size.

The total social costs associated with the gap between pooling and separating equilibria are also determined by the proportion of the population at each type.¹⁴⁹ In Figure 2, there are also two possible distributions of types, f_1 and f_2 , with the latter being more dispersed, representing a more heterogeneous population. If most of the population is located near the mean type (distribution f_1), the welfare losses from privacy associated with both x^*_L and x^*_H will be relatively small, as for the vast majority of the affected population the gap between optimal action and the uniform action taken in a pooling equilibrium is relatively small. Thus, even if there are large gains from separation, when most of the population is homogenous over the trait of interest, the inefficiencies from pooling will be small. Alternatively, if the population varies a lot over the trait in interest (distribution f_2), total social losses will be larger for both x^*_L and x^*_H . Because a much smaller proportion of the population is centrally located, there will be non-trivial social losses for pooling even if the gains from separation are relatively small.

¹⁴⁹ Formally, social losses are: $\int_{\underline{R}}^{\bar{R}} [V(x^*(R_i)) - V(\bar{x}(\bar{R}))] f(R) dR$.

FIGURE 2. LOSS FROM MISMATCH OF OPTIMAL ACTION

As an illustration, suppose that R measures ability successfully to complete law school and x is the discount on tuition. If there are large gains from matching types to tuition, $x^*(R)$ will be relatively steep; those who are unlikely to succeed (low R s) should be discouraged with high tuition from attending and wasting their time and money, whereas those with high aptitudes for a legal career should be given large discounts to encourage them 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 absolute cost associated with a particular mismatch may be large, if the pool of applicants is relatively homogenous (f_1), 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 (f_2), the losses from the most severe mismatches receive more weight.

2. Identifying Optimal Restrictions on Big Data

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:¹⁵⁰

$$P(z)*H + \theta z.$$

$P(\bullet)$ is the probability of an accident, z measures effort to avoid an accident, H is the harmed caused by an accident, and θ is the marginal cost of care.

Adapting this model to the case of privacy harms from big data classification, we can think of z as the decision to limit big data classifications. For example, low levels of z may be simple notice requirements, with z increasing as restrictions ratchet up from opt-in requirements, to use restrictions, and finally collection restrictions, like data minimization requirements.¹⁵¹ Higher levels of z , therefore, correspond to lower likelihoods that a consumer will suffer an intrinsic privacy harm, H . 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

¹⁵⁰ See Shavell, *supra* note 134, at 177–79.

¹⁵¹ See Ramirez, *supra* note 56 (“[i]nformation that is not collected in the first place can’t be misused”); Alvaro Bedoya & David Vladeck, Center for Privacy & Technology at Georgetown Law Center, *Comments On Big Data and Consumer Privacy in the Internet Economy*, (Aug. 5, 2014).

data, at least some of which are likely to be beneficial.¹⁵² These costs are represented by θ .

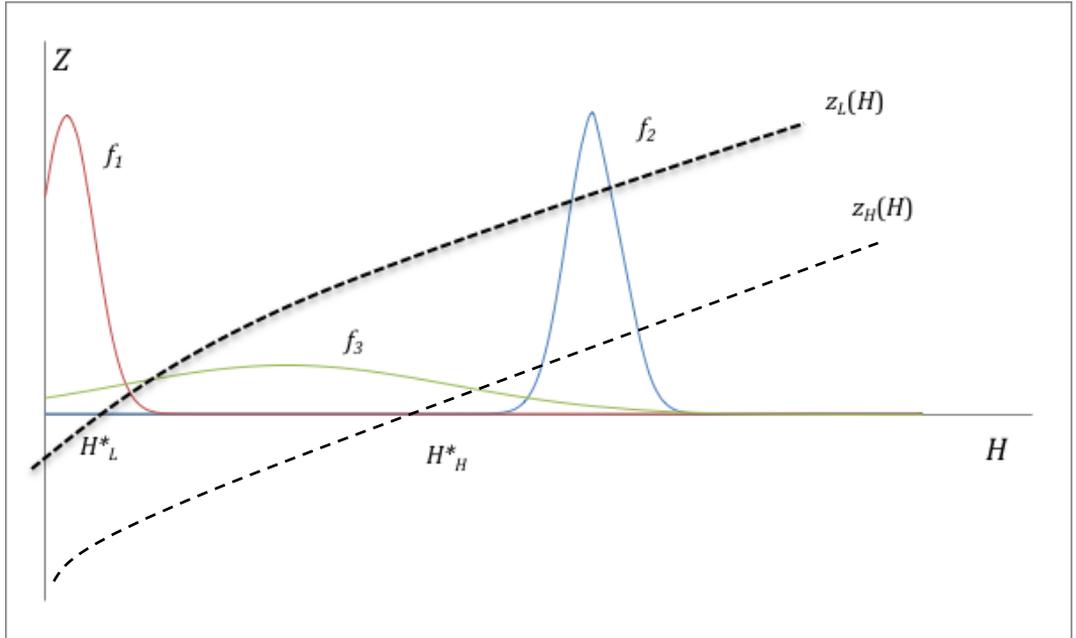
In Figure 3, there are two curves, z_L and z_H , which map optimal levels of big data use for various levels of harm—clearly, the higher the harm, the more restrictions are placed on big data (*i.e.*, higher z).¹⁵³ The differences between these curves are the benefits from big data-driven separation.¹⁵⁴ For z_L , the benefits from separation are small, so retarding collection or use of data comes at little cost (θ_L). The marginal 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^*_H , is met.

FIGURE 3. OPTIMAL BIG DATA RESTRICTIONS WITH HETEROGENEOUS HARMS AND BENEFITS

¹⁵² 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.” PRESIDENT’S COUNCIL OF ADVISORS ON SCI. & TECH., REPORT TO THE PRESIDENT: BIG DATA AND PRIVACY: A TECHNOLOGICAL PERSPECTIVE 49 (May 2014). The PCAST Report notes the potentially crippling expenses associated with enforceable collection regulation. *Id.* 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.”).

¹⁵³ The curves are concave because of the diminishing marginal effectiveness of additional precautions: $\frac{\partial z^*}{\partial H} = -\frac{P'}{HP''}$.

¹⁵⁴ It can be shown that $\frac{\partial z^*}{\partial \theta} = -\frac{1}{HP''} < 0$.



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 of 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 use the mean to approximate the level of harm for most of the population.

These distributions underscore how crucial it is to have information about intrinsic privacy harms when calibrating policy with respect to the use of big data to classify individuals. Absent information on intrinsic harms, restricting big data merely because it

uncovers personal attributes that will be used for classification (e.g., prices, insurance rates, or credit terms) risks depriving consumers and society of the benefits from separation by confusing benefits for harm. 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, credit worthiness, driving ability, taste in clothes or food, or sexual orientation. 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.

3. Measuring Intrinsic Privacy Harm

What do we know about intrinsic privacy value? The short answer is not much. The available empirical evidence on how consumers feel about observation could shed some light on potential intrinsic harms from classification, but it provides little guidance. Survey data show that consumers care about privacy, yet revealed preferences suggest stated concerns may be exaggerated. In a recent Pew Poll, 65% of respondents say that “controlling what information is collected about you” is “very important.”¹⁵⁵ At the same time, consumers increasingly participate in online activities that reveal personal data to known and unknown third parties; the percentage of online adults engaging in social media rose from 8% in 2005 to 72% in 2013.¹⁵⁶ Additionally, the health tracking market has exploded in recent years.¹⁵⁷ Further, very few people bother to opt-out of online tracking or adopt privacy-protecting technology, like searching via Duck, Duck, Go! or using the TOR browser.¹⁵⁸ Indeed, Acquisti,

¹⁵⁵Mary Madden & Lee Rainie, *Americans' Attitudes About Privacy, Security, and Surveillance*, PEW RESEARCH CENTER (May 15, 2015), <http://www.pewinternet.org/2015/05/20/americans-attitudes-about-privacy-security-and-surveillance/>.

¹⁵⁶ Joanna Brenner & Aaron Smith, *72% of Online Adults are Social Networking Site Users*, PEW RESEARCH CENTER (Aug. 5, 2013), <http://www.pewinternet.org/2013/08/05/72-of-online-adults-are-social-networking-site-users/>.

¹⁵⁷ 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/>.

¹⁵⁸ See Maurice E. Stucke & Allen P. Grunes, *No Mistake About It: The Important Role of Antitrust in the Era of Big Data*, ANTITRUST SOURCE (April 2015), http://app.antitrustsource.com/antitrustsource/april_2015?pg=25#pg25.

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.¹⁵⁹ Thus, although consumers are concerned about being observed, their 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.

Some researchers have attempted to measure intrinsic privacy valuation, but the thin extant literature provides little that is generalizable. 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. Consistent with an endowment effect of privacy, the authors find that consumers appear to value privacy more when they are asked to sell it than when they must purchase it.¹⁶⁰ Consumers' willingness to divulge private information also appears to depend on context and cues.¹⁶¹ Further, perceptions about the ability to control one's information impact willingness to share personal information.¹⁶² Other researchers have found that consumers would be willing to accept small discounts and purchase recommendations in exchange for personal data¹⁶³ and that they exhibit low willingness to pay for protection from telemarketers.¹⁶⁴ For example, one study finds that consumers are willing to pay an additional one to four dollars for a

¹⁵⁹ See Alessandro Acquisti et al., *The Economics of Privacy*, J. ECON. LIT. (forthcoming 2016) (manuscript at 37-38), available at http://people.duke.edu/~crtaylor/Privacy_Survey.pdf.

¹⁶⁰ See Alessandro Acquisti, Leslie K. John, & George Lowenstein, *What is Privacy Worth?*, 42 J. LEG. STUD. 249, 254 (2013).

¹⁶¹ See Alessandro Acquisti, Leslie K. John, & George Lowenstein, *Strangers on a Plane: Context-Dependent Willingness to Divulge Sensitive Information*, 37 J. CONSUMER RES. 858, 859 (2011).

¹⁶² See Laura Brandimarte, Alessandro Acquisti, & George Lowenstein, *Misplaced Confidences: Privacy and the Control Paradox*, 4 SOC. PSYCHOL. & PERSONALITY SCI. 340, 345 (2012).

¹⁶³ See Dan Cvreck, Marek Kumpost, Vashek Matyas, & George Danezis, *A Study on the Value of Location Privacy*, 5 ACM WORKSHOP ON PRIVACY IN THE ELECTRONIC SOC'Y 1 (2006). For a review of the empirical literature, see also Acquisti et al., *supra* note 160, at 39.

¹⁶⁴ 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), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1000533.

hypothetical smartphone app that conceals location, contacts, text content, or browser history from third-party collectors.¹⁶⁵

In highly relevant research, recent work by Benjamin Wittes and Jodie Liu suggests that people are more concerned with privacy intrusion from proximate observation by individuals than distant observation by computers.¹⁶⁶ 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.¹⁶⁷ This research suggests that intrinsic privacy harms from big data predictions may be overstated to the extent that these predictions are made, and known, only by algorithms.

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. In light of the costs associated with overdeterrence of beneficial practices, policymakers should proceed with caution. When policymakers 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

¹⁶⁵ Scott Savage & Donald M. Waldman, *The Value of Online Privacy* (2013), http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2341311.

¹⁶⁶ Benjamin Wittes & Jodie Liu, *The Privacy Paradox: The Privacy Benefits of Privacy Threats*, CENTER FOR TECHNOLOGY INNOVATION AT BROOKINGS 9-10 (May 2015), http://www.brookings.edu/~media/research/files/papers/2015/05/21-privacy-paradox-wittes-liu/wittes-and-liu_privacy-paradox_v10.pdf.

¹⁶⁷ See also Stephanie Mathson & Jeffrey Hancks, *Privacy Please? A Comparison Between Self-Checkout and Book Checkout Desk for LGBT and Other Books*, 4 J. ACCESS SERVICES 27, 28 (2007).

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.¹⁶⁸ In the context of big data, “too much care” can be mean self-limiting beneficial uses of data.

4. Developing Defaults

In an ideal world, restrictions on the use of big data to make predictions would be tailored to idiosyncratic sensitivities. But in the real world, policy choices are lumpy rather than continuous; in most scenarios, a homogenous rule must be applied to a heterogeneous population. Thus, the task becomes one of deciding a regulatory default. One could imagine three broad categories of regulation: first, “permissionless innovation” would place the burden on those advocating restrictions on big data to show harm;¹⁶⁹ second, there could be disclosure requirements with either opt-out or opt-in consent; and third, there could be presumption in favor of restricting big data—with respect to use or data collection—with the burden on those opposing restrictions to show that the benefits outweigh the privacy harm.¹⁷⁰ In this way, rather than relying on one-size-fits-all *ex ante* regulation proposed by some scholars,¹⁷¹ policy makers can tailor regulatory stances to the situation.

The framework developed in this Article can help identify factors that should influence the proper default regulatory posture, and it has the advantage of relying on information that is more readily available than that on intrinsic privacy values. For example, the

¹⁶⁸ 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 134, at 224–29.

¹⁶⁹ See Adam Thierer, *Permissionless Innovation: The Continuing Case for Comprehensive Technological Freedom*, MERCATUS CENTER 9 (2015).

¹⁷⁰ See Ramirez, *supra* note 56 (“[i]nformation that is not collected in the first place can’t be misused.”).

¹⁷¹ See *supra* notes 51–53 and accompanying text.

distribution of creditworthiness is generally known,¹⁷² 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 help map out $z^*(H)$, which will provide a threshold of harm that would be necessary to justify big data restrictions. Further, although intrinsic privacy harm distributions are unknown, there is rough agreement on the sensitivity of certain types of information. Health information, for example, especially concerning conditions that carry potential stigma (e.g., mental health conditions or sexually transmitted disease), is widely considered sensitive—so much so that we have developed an entire statutory scheme to guard its confidentiality.¹⁷³ Similarly so for information concerning children.¹⁷⁴ On the other hand, there is likely to be wide agreement that information about driving or eating habits is not the kind of information that most consider sensitive. Thus, in many cases we can at least have some rough sense of whether we are confronting a distribution that looks more like f_1 or f_2 . By marrying this estimation with what we know about the likely benefits from promoting separation in a market, we can arrive at an informed default posture.

i. Default in Favor of Big Data Classification

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

1. Big data is aimed at separation on a trait that is relatively widely dispersed;
2. Gains from reducing adverse selection and moral hazard are likely to be large; and
3. Predictions involve non-sensitive traits.

This scenario is represented by an optimal restriction curve z^*_H in Figure 3. When the gains from big data are large (optimal restriction curve z^*_H) and the harms are small (distribution f_1), the

¹⁷² See, e.g., *Search Results for: Distribution*, FICO BLOG (2017), <http://www.fico.com/en/blogs/?s=distribution>.

¹⁷³ See Health Insurance Portability and Accountability Act, 42 U.S.C. § 1320d (2010). Note that while HIPAA considers mental health information sensitive information, it does not distinguish mental health information as requiring more protection than other medical information.

¹⁷⁴ See Children's Online Privacy Protection Act, 15 U.S.C § 6501 et seq. (1998).

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 restrictions on big data. The overwhelming demand for concealment in this scenario is strategic. It is easy to see that adopting a permissionless innovation stance would maximize welfare because adopting even a small level of regulation would impose a cost on all but the most privacy sensitive. Accordingly, the presumption in these cases should be to allow the use of big data unless the trait in question is one over which discrimination laws forbid classification.

So, what types of practices likely fall into this bin? The most obvious transactions where collection should be favored are those in which 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 postings 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. Additionally, the predictions here are geared toward identifying *different* risk profiles in an attempt to expand insurance coverage or access to credit, rather than discerning reservation prices for *similar* risk profiles to gain a larger share of the surplus. Thus, this use of big data is unlikely to be dissipative.

Analytics used for online and offline marketing are another candidate. First, the gains from identifying distinct tastes and preferences are not likely to be trivial.¹⁷⁵ Although the demand for concealment in this case is not strategic, most people are unlikely to suffer any substantial privacy harm—the observation and analysis of online and offline shopping habits by an anonymous algorithm to

¹⁷⁵ See Avi Goldfarb & Catherine E. Tucker, *Privacy Regulation and Online Advertising*, 57 *MGM³T SCI.* 57 (2011); J. Howard Beales, III, *The Value of Targeted Advertising* (2010), http://www.networkadvertising.org/pdfs/Beales_NAI_Study.pdf.

make predictions about the types of goods and services one likes involves the collection of, and classification relating to, relatively non-sensitive information.¹⁷⁶ Further, being classified as enjoying one type of good over another carries no stigma in most circumstances—especially because the one making the prediction is an anonymous server and the only one seeing the prediction in most cases is the consumer. If the concern is over people being excluded from certain goods or services erroneously (*e.g.*, funneled into only subprime offers), this is not a privacy harm but rather the domain of antidiscrimination law.¹⁷⁷ For example, the Equal Credit Opportunity Act¹⁷⁸ and its implementing regulations¹⁷⁹ govern disparate impact or treatment of protected groups that may arise from big data analytics.¹⁸⁰

ii. Default Against Big Data Classifications

The following characteristics militate toward a more aggressive regulatory stance:

1. Homogeneous populations over the trait in question;
2. Little problem with adverse selection or moral hazard; and
3. Predictions involving highly sensitive traits.

For example, if privacy harm can be represented by distribution 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). Here, if one set the level of restriction to the optimal level for the mean of the distribution, some of the population would suffer from too many restrictions, while others would suffer from too few. Nonetheless, because the distribution is tightly clustered around the mean, and even the least privacy sensitive member of this population demands some privacy regulation, allowing controls on big data here maximizes welfare. When considering *ex ante* restrictions, however, policymakers should heed the advice of the President's

¹⁷⁶ See, *e.g.*, *Gaos v. Google, Inc.*, 2012 WL 1094646 at *2 (Mar. 29, 2012) (finding that alleged harm from Google's "dissemination of Plaintiff's search queries to third parties" was insufficient for Article III standing).

¹⁷⁷ Further, errors in classification are likely to be corrected by competition. See FED. TRADE COMM'N, *supra* note 61, at 18–19.

¹⁷⁸ 15 U.S.C. §§ 1691 *et seq.* (2014).

¹⁷⁹ 12 C.F.R. § 1002 (2011).

¹⁸⁰ See, *e.g.*, FED. TRADE COMM'N, *supra* note 61, at 18-21.

Council of Advisors on Science and Technology, and they should be hesitant to prohibit collection of data.¹⁸¹ Such regulation 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, at least some of which are likely to be beneficial.

As explained above, distributions like f_2 may represent privacy harms involving sensitive health information or information about children. Rather than complete prohibition of information collection, laws like HIPAA¹⁸² and COPPA¹⁸³ require opt-in consent. The social benefits from regulation are even more clear when the benefits from big data are small. For example, this situation 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 given the near homogeneity of the population with respect to this trait. Further, allowing these predictions to be used could discourage discovery of this information in the first place, which could be useful in the hands of the sufferer. Indeed, Congress appears to have made this determination when enacting the Genetic Information Nondiscrimination Act.¹⁸⁴ Notably, this law focusing on restrictions on use, not collection, in recognition of the potential social value that can be derived from genetic information.

Similarly, using big data to create “sucker lists” is an example of expending resources merely to determine a class of people from

¹⁸¹ 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.” PRESIDENT’S COUNCIL OF ADVISORS ON SCI. & TECH., *supra* note 152, at 49. The PCAST Report notes the potentially crippling expenses associated with enforceable collection regulation. *Id.* 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.”).

¹⁸² 45 C.F.R. § 164.500 et seq. (2013).

¹⁸³ 16 C.F.R §§ 312.4–312.5 (2013).

¹⁸⁴ 42 U.S.C. §§ 2000ff et seq. (2008). Further, several states have similar provisions. See e.g., Amalia R. Miller & Catherine Tucker, *Privacy Protection, Personalized Medicine, and Genetic Testing* (2015),

http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2411230.

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.¹⁸⁵ A default in favor of big data restrictions in these cases would make sense. For example, the FTC could use its unfairness authority to address big data uses to create sucker lists under the assumption that they are likely to create substantial harm and there are no countervailing benefits to consumers or competition.¹⁸⁶

iii. The Hard Cases

Of course, the hard cases occur when there is no agreement on the sensitivity of what big data will reveal. When the distribution looks like f_3 , for example, the average level of harm provides little information on how most people value privacy because almost all of the population suffers harm away from the mean. Absent an accurate estimate of harms, we should instead focus on 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 a notice with an opportunity to opt-out for the specific use.

The Target incident may fit into this category. The distribution of intrinsic privacy harm from having one's pregnancy predicted by an algorithm is likely to look like f_3 . Although the predictions have to do with marketing, because privacy is concerned, more people are more likely to suffer some privacy harm than behavioral targeting on non-sensitive traits represented by f_1 . Indeed, the uproar this incident caused in the privacy activist community suggests that some segment of the population views these facts as giving rise to non-trivial harm. At the same time, it is a dubious proposition that second-trimester pregnancy status rises to the level of medical conditions that come with stigma or embarrassment depicted in f_2 . Although some women

¹⁸⁵ See Vladeck, *supra* note 148, at 162.

¹⁸⁶ 15 U.S.C. § 45(n) (2006). At the same time, the FTC should use caution in defining a "sucker list" so as not to deprive consumers of legitimate offers.

may want to conceal their pregnancy (e.g., from disapproving parents or from current or prospective or current employers¹⁸⁷), this situation is unlikely the case for most women.¹⁸⁸ Further, even if you want to conceal your pregnancy from the world, a Target algorithm correctly predicting your pregnancy is a far cry from having it revealed to the world.¹⁸⁹ 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 likely to be slim.

Since there is liable to be a large range of intrinsic harm associated with something like the Target incident, we should turn to the benefits. Target used data from its baby shower registry—which provided it with a list of women with known due date—to analyze shopping habits, with the goal of being able to send offers to women in their second trimester.¹⁹⁰ 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.¹⁹¹ 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.¹⁹² Thus, there are clear benefits to

¹⁸⁷ Alissa Quart, *Why Women Hide Their Pregnancies*, N. Y. TIMES, Oct. 7, 2012.

¹⁸⁸ 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, (2017), <http://www.webmd.com/baby/features/pregnancy-stages-baby-body>.

¹⁸⁹ *See* Wittes & Lui, *supra* note 166.

¹⁹⁰ Charles Duhigg, *How Companies Learn Your Secrets*, N. Y. TIMES MAG. Feb. 16, 2012.

¹⁹¹ 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:

ΔThese gains are magnified if the targeted goods were more price elastic. Total s

¹⁹² Even without price reductions, mere advertisements are likely to increase demand for prenatal vitamins and hence total surplus. *See, e.g.*, Dhaval Dave & Henry Saffer, *Impact of Direct-to-Consumer Advertising on Pharmaceutical Price and Demand*, 79 S. ECON. J. 97 (2012) (finding that direct-to-consumer broadcast advertising of prescription drugs increased demand by about 12%); Pauline M. Ippolito & Alan D. Mathios, *Information, Advertising and Health Choices: A Study of the Cereal Market*, 21 RAND J. ECON. 459 (1990) (finding that removal of advertising ban on advertising health benefits of fiber increased consumption of fiber). Note that these studies are not specific to prenatal vitamins; they investigate

separation here that would be lost if stringent *ex ante* regulations (e.g., a privacy review board) kept Target from making predictions about pregnancy or even collecting the data in the first place.¹⁹³ Thus, prohibiting stores from collecting such data or using it in algorithms to predict pregnancy would be an inappropriate regulatory response. Rather than *ex ante* limitations, requiring disclosure of such data uses and allowing consumers the ability to opt out from marketing would be a more sensible approach. This approach will allow consumers to self-select based on their privacy sensitivity.¹⁹⁴ Here, the FTC could use its power under Section 5 to enforce promises with respect to collection and use of data.¹⁹⁵

CONCLUSION

Big data may raise genuine privacy concerns, but anxieties over separation by itself risk confusing benefits for costs. Calls to regulate big data must be careful to distinguish between productive and dissipative concealment of private facts. Restricting the use of data to draw inferences merely because they will cause some people to suffer worse terms is unambiguously welfare-reducing, and it is unfair: it forces those with relatively good attributes to subsidize both those with below average attributes and those with extreme preferences for privacy. Further, in many cases the relatively good types in disadvantaged populations will be the ones who pay the most for forced pooling. Preventing firms from using big data to discover these people only condemns them to suffer higher prices, less credit, and fewer job opportunities.

the relationship between direct to consumer advertising of products in the medical field and health claims in breakfast cereals.

¹⁹³ See Bedoya & Vladeck, *supra* note 151 (“strong *ex ante* use limitations could have stopped Target from identifying pregnant women through their purchases.”).

¹⁹⁴ See Howard Beales et al., *The Efficient Regulation of Consumer Information*, 24 J.L. & Econ. 491, 513 (1981) (“informational remedies allow consumers to protect themselves according to personal preferences rather than place on regulators the difficult task of compromising diverse preferences with a common standard”).

¹⁹⁵ See, e.g., *FTC v. CompuCredit*, No. 1:08-cv-1976-BBM-RGV (N.D. Ga. 2008) (alleging deception when, *inter alia*, credit card marketing company failed to disclose that behavioral credit scoring would lead to the reduction in credit lines if credit cards were used to pay for certain things).

At bottom, policies that prevent the discovery of information to protect strategic privacy are not animated by a desire to preserve dignity or autonomy, but rather to guard against different outcomes based on *reactions to that information*. Curtailing the use of big data in the name of privacy is the wrong approach here. Instead, it is far better to deal with this separation anxiety directly, through antidiscrimination law, which embodies society's choices about the dimensions over which separation is forbidden.