

Consumer Protection in Cyberspace

Oscar H. Gandy, Jr.

*University of Pennsylvania
Annenberg School for Communication
<http://www.asc.upenn.edu/usr/ogandy>
ogandy@asc.upenn.edu*

Abstract: *This critical essay introduces the problem of discrimination enabled through the use of transaction-generated-information derived from the analysis of user behaviors within the network environment. The essay begins by describing how segments of the population that are already vulnerable become further victimized through the strategic use of discriminatory algorithms in support of identification, classification, segmentation, and targeting. In response, it evaluates a set of policy options that might be used to limit the harm and compensate the victims of these inherently dangerous technologies. Traditional approaches that stress the protection of privacy through restrictions on the collection and use of personal information are compared with alternatives based on individual and class actions under tort law, as well as more traditional regulatory approaches developed in the area of consumer products safety and environmental regulation.*

Keywords: Discrimination, surveillance, privacy, online, regulation of technology, data protection, measurement error, personal information, harm, liability, environmental regulation

This critical essay is an attempt to shift the terms of debate within the community of scholars concerned with privacy regarding the potential for regulatory intervention to condition its survival (Bennett 2011). Although privacy, and surveillance, and their weaker cousin, data protection, are the dominant frames through which the myriad problems associated with the collection and use of personal and transaction-generated-information (TGI) tends to be discussed, I want to propose an alternative. Although my ultimate goal is an engagement with the conjoined problems of discrimination and inequality, my chosen point of entry is an assessment of information technology (Gandy, 2009; 2010). In this case, the technology is the family of statistical resources within the analysts' toolkit that transform TGI into actionable intelligence (Davenport & Harris, 2007).

The essay begins by describing how already burdened segments of the population become further victimized through the strategic use of sophisticated algorithms in support of the identification, classification, segmentation, and targeting of individuals as members of analytically constructed groups. Although the use of these decision support systems is routinely justified in terms of their contributions to corporate profitability and the management of risk, there is growing recognition of the negative externalities, or harms that are generated at the same time (Allhoff, 2009; MacCarthy, 2010). A brief review of the broad scope of these harms will be provided.

We then consider a set of policy options that might be pursued in an attempt to limit the harm and compensate the victims of these inherently dangerous technologies. Traditional approaches that stress the protection of privacy through restrictions on the collection and use of personal information will be compared with alternatives based on individual and class actions under tort law, as well as more traditional regulatory approaches developed in the area of consumer products safety and environmental regulation.

1. Statistical Discrimination

I characterize these systems, devices and procedures as discriminatory technologies, because discrimination is what they are designed to do (Davenport & Harris, 2007). Their value to users is based on their ability to sort things into categories and classes that take advantage of similarities and differences that seem to matter for the decisions users feel compelled to make (Gandy, 1993).

While developing actionable intelligence in support of strategic discrimination is the purpose for which these sophisticated analytics are brought into use (Gandy, 2009, pp. 19-34), this process ultimately depends upon the successful identification, classification, and evaluation of people, places and things as targets of interest (Wilkinson, 2008). Although it is convenient to distinguish between these different intermediate goals by suggesting that identification is about determining who you are, while classification is about determining what you are, there are additional distinctions to be drawn. The most important distinctions are those upon which economic valuations are made.

All of these assessments act as aids to discrimination - guiding a choice between entities. Each choice represents an action that affects the status of entities as winners or losers in a game of chance (Dahrendorf, 1979). In many cases, the decisions made by the users of sophisticated analytics determine the provision, denial, enhancement, or restriction of the opportunities that citizens and consumers face both inside and outside formal markets. The application of analytics is migrating rapidly from the realm of face-to-face interactions to the realm of computer-mediated-interactions.

What we are concerned about here is a technologically enhanced process that is endlessly repeated across the terrain of cyberspace that Mark Andrejevic (2007) has characterized as a "digital enclosure". The statistical discrimination enabled by sophisticated analytics contributes to the cumulative disadvantage (Gandy, 2009) that weighs down, isolates, excludes, and ultimately widens the gaps between those at the top, and nearly everyone else (Wilkinson & Pickett, 2010). Although observers have tended to focus on the use of these systems in support of targeted advertising online, their reach is far more extensive. It covers access to a range of goods and services, including markets for finance and housing, as well as health care, education, and social services. Its reach has also expanded rather dramatically into components of the criminal justice system in ways that threaten the delivery of other services through government bureaucracies (Gandy, 2009).

Although we tend to think about the users of discriminatory technologies as rational beings acting in pursuit of their interests (as they understand them), it is important to point out that increasingly these "choices" are being made by automated digital agents completely without assistance or interference from human beings (Tokson, 2011).

For the purposes of this paper, I am setting aside discriminatory acts that are based on invidious distinctions reinforced by animus and malicious intent. Instead, I emphasize the way that statistical discrimination compounds the disadvantages that the structural constraints we readily associate with race, class, gender and cultural identity influence the life chances that shape the opportunity sets that people encounter. This is what cumulative disadvantage means in practical terms (Gandy, 2009, pp. 74-76).

My challenge then becomes one of demonstrating that the disparate impacts that we readily observe are the result of accidental errors, or other unintended byproducts or externalities.

2. Errors and Externalities

Errors are not only possible, but they are likely to occur at each stage in the process of assessment that proceeds from identification to its conclusion in a discriminatory act. Error is inherent in the nature of the processes through which reality is represented as digitally encoded data (Hayes, 1993). Some of these errors will be random, but most will reflect the biases inherent in the theories, and the goals, the instruments and the institutions that govern the collections of data in the first place (Hobart & Schiffman, 1998).

Because the decisions of interest to us here are those that are made on the basis of assessments of groups, we have to consider the myriad sources of error that are encountered along the way. The first, and perhaps most important source of error is in the nature of the samples that are drawn from the population of interest (Hand, 2006).

Sampling error is unavoidable, even when those samples are drawn by well crafted, and dutifully followed rules. The factors that combine to determine the nature and scope of these errors are vast (Anderson & Sclove, 1986). The errors that are common to convenience samples, or other collec-

tions of data that ignore “best practices” are not only substantial, but they defy reasonable estimation. Additional errors are introduced into the mix at the moment of measurement, or capture, as well as during the transformation of first impressions into digital forms (Agre, 1997).

Although researchers and statisticians have made progress in identifying the sources of error in our attempts to characterize the past, we are just beginning to understand how much more likely there are to be errors in our predictions about how things will be in the future (Gigerenzer et al., 1989). Because these predictions are based on data and theories about occurrences in the past, they will necessarily be limited by a lack of information about and understanding of the conditions, circumstances, and expectations that will govern decisions that will be taken at different points in the relevant future (Giddens, 1984).

In addition to the harms that are introduced through error, there are additional sources of harm that need to be brought into the mix. I refer here to what economists discuss under the heading of externalities. They are referred to as external effects, or spillovers, because they affect actors, entities, or third parties that are not part of the primary economic transaction (Baker, 2002).

Markets are said to depart from optimal performance because producers and consumers can neither capture the benefits, nor be compelled to pay the costs of the externalities that affect third parties. In addition, because the attention of market actors is focused on the pursuit of their own interests, they tend to ignore many critical changes in the environment that have been generated as byproducts of their actions (Tenner, 1997; 2004). They pay even less attention to the distribution of those harms, and how they tend to cumulate as barriers to opportunity for the already disadvantaged (Gandy, 2009). Regulatory intervention by public and private agencies is generally relied upon to generate a second best, or suboptimal result.

3. The Policy Response

Because the errors and the harms that accompany the rapidly expanding deployment of analytics into the networked environment (Ayres, 2007a) are both substantial and extensive, as well as being maldistributed, many of us have been seeking answers and alternatives within the sphere of law and regulatory policy. We will begin here by examining some of the problems and prospects that arise when we place our hopes on what Colin Bennett (2011, p. 486) has characterized with regard to privacy as a “regime of governance and as a set of practices.”

4. Privacy and its discontents

Privacy, like participatory democracy, is more of an ideal than a fact of life. Privacy is an ideal marked by considerable variation in what it actually means for most of us. Most of our efforts to protect privacy have been developed under the banner of data protection.

We have tended to focus on data as the thing to be protected, and only indirectly on privacy as the state of being that might be enjoyed as a result of having protected those data. Unfortunately, as a result of this historical focus, most of the laws that protect data have been distorted by constructions of data as property. They have been constrained further by a tendency to think about this intangible property as commodities that can be bought, sold, traded and conveyed through markets.

4.1. Personally Identifying Information

While a focus on data as property is already troublesome, it pales in comparison with the problems that are associated with specifying just which data deserve special attention because of its status as “personally identifying information” or PI.

David Phillips (2004) has invited us to think about three different kinds of PI that relate in part to its creation, and in part to its primary function - that of identification. According to Phillips, lexical identification establishes a link between an entity and its name. Indexical identification differs to the extent that one can point to, or as we might say, “reach out and touch” a particular individual. This

becomes especially relevant in light of Phillip's analyses (2003; 2009) that focus on mobility and the forms of identification that incorporate information about a person's physical location at some specific point in time. In Phillip's view (2004), descriptive identification involves the assignment of attributes to a particular entity or class that helps to define them in relation to differentially identified others.

What is particularly troublesome about this approach to PI is the fact that our tendency to focus on the identification of individuals keeps us from understanding how those individuals are affected by the descriptive and indexical identification of the groups to which they are said to belong.

This is not only, or even primarily a concern about the kinds of groups that have already achieved some level of protection within democratic societies, perhaps under the umbrella of civil and human rights. Our concern here is about the kinds of "groups" that tend to be defined rather idiosyncratically by the organizations that bring them into being through statistical analysis.

On occasion, the names applied to these groups become more widely known through a process of social osmosis or leakage across discursive boundaries. But even those groups tend to remain politically powerless, unless someone takes the initiative to form an organization in order to pursue identifiable common interests. One important example would be the American Association of Retired Persons (AARP) that has lobbied successfully over time in support of the interests of senior citizens in the US.

4.2. Fair Information Practices

Actually, by focusing our approach to privacy through the lens of data protection, we have limited the ability of Fair Information Practices (FIPs) to do the work they were initially designed to do (Solove & Rotenberg, 2003). While we've seen the number of principles expand from five to eight and more over time, we've also seen some rather severe contractions in this range, as with the US Federal Trade Commission (FTC) emphasizing something called "notice and choice" (FTC, 2010, p. 20), while ignoring more fundamental commitments to limiting the collection, use, and sharing of personal information.

Naturally, privacy advocates, especially those in the United States, note that these FIPs are principles, rather than regulations. They are therefore quite limited and weak with regard to oversight, accountability and enforcement. For many, the defense of these principles is seen to chronically lag behind changes in the nature of information technology and marketplace standards with regard to the use of personal data and TGI (Gellman, 2008).

There are also substantial concerns being expressed about the usefulness of principles that suggest individuals should have the right of access and correction for data about themselves (Gomez, Pinnick, & Saltani, 2009). Certainly, the right to challenge a factual statement about whether a purchase was made, or whether a bill was paid, represents a meaningful opportunity for the consumer. But how does a consumer begin to challenge the accuracy of a credit score, or more critically, a prediction regarding the probability of default, or some other determination of creditworthiness that has been based on some complex multivariate and proprietary assessment tool?

Similar concerns arise with regard to meaning of informed consent, especially as it relates to the subsequent uses to which transaction-generated-information, or TGI can be put (MacCarthy, 2010). Truly informed consent cannot exist because of the difficulty that all of us would face in trying to predict how any particular bit of information might be used to shape the opportunities and constraints we might face at some point in the future.

Considering that it is virtually impossible to make use of the internet without generating TGI, it is disingenuous, at best, to suggest that individuals could actually provide informed consent to the capture, collection, analysis and subsequent use of this information.

We also need to consider the kinds of information networks that are implicated when we think we are merely referring to ourselves, such as when we update our profiles on some social networking site (Gelman, 2009; McGeeveran, 2009), or when we answer a public opinion or marketing survey. It is important that we realize that the people in the so-called "representative samples" utilized

by pollsters are actually providing strategic information about people who have not granted their consent, in part because they have not been asked. It is in the nature of so-called representative samples that the voluntarism of the compliant actually places the rest of us at risk (MacCarthy, 2010).

The most important, and therefore the most assiduously ignored principle is the one that is concerned with limiting the uses to which personal data can be put (Solove & Rotenberg, 2003, pp. 724-726). The problems of subsequent use become especially salient with regard to the class of activities that proceed under the heading of data mining-- a term that refers to a variety of applied statistical techniques that are used to derive meaning from the patterns that emerge through the processing of massive amounts of data (Gandy, 2010).

4.3. Data Mining

By its very definition data mining involves subsequent uses of information. And it seems unlikely that many of us could succeed in demonstrating in some legally meaningful way, that those uses were "incompatible" with the original purposes for their collection (Solove, 2008).

There are ongoing debates about the kinds of restrictions we might establish on the incorporation of data from public and private sources into the profiles that describe individuals (FIDIS, 2009a). With the rapid growth of cloud computing and the storage of all sorts of data on remote computers, our notions of what a reasonable expectation of privacy should be, clearly have to be reassessed (Wright et al., 2009).

Data mining and secondary uses of information is an especially salient concern with regard to the "Googlization" of virtually everything we might choose to do in a networked environment (Zimmer, 2008). In addition, the fact that social networking sites are routinely indexed by their corporate parents means that massive amounts of PI, including some incredibly sensitive materials, are already likely to be part of some data mining operation (Tokson, 2011, pp 46-48). As Helen Nissenbaum notes (2010, p. 63): "MySpace, for example, is developing a method to harvest user profiles for demographic information that can be provided to advertisers, enabling these advertisers to target their audiences precisely," or at least precisely enough for the current state of the market for consumer profiles.

The same kinds applied statistical power are also likely to eliminate the protections we once thought we had in the promises made by database managers to "anonymize" or "de-identify" their datasets. On the one hand, we see data mining activities generating new aggregations, and then characterizing, identifying, and evaluating their "signatures" or digital footprints. On the other hand we find another group of "wizards" making reidentification an almost trivial pursuit (Ohm, 2010).

We need to be reminded that privacy interests, as the courts usually interpret them, apply only to the exposure of PI to other human beings. Clearly, we're not supposed to be embarrassed if some computer at the telephone company takes note of what we've been looking at lately (Tokson, 2011). But of course, embarrassment is not primarily what we're concerned about here. Our concern is about discrimination and cumulative disadvantage, outcomes that are being generated on a continual basis by autonomous digital agents across the network (Hildebrandt, 2008).

5. The Problem with Torts

Given the importance of harm and its sources to the protection of privacy, let us turn our attention to an area of law in which assigning liability for harm is a central concern (Twerski & Henderson, 2009; Vidmar & Wolfe, 2009; Polinsky & Schavell, 2010).

Tort law has developed in part as a function of its primary use, and the nature of the resources that contending parties can bring to bear on the considerations of the courts (Rubin & Shepard, 2008). We know that defendants are fairly likely to succeed in avoiding liability in medical malpractice cases (Kritzer, 2007; Mello et al., 2007), while plaintiffs are relatively more successful in other tort domains (Nockleby, 2007).

Limitations on contingency fees, and damage awards weigh especially heavy on individual plaintiffs who are pursuing compensation for injuries that are both slight, and difficult to prove (Sandefur, 2008). As a result, lawyer-advocates willing to risk their own limited funds on these cases are relatively few in number. At the same time, in the US at least, increasing limitations on the nature of mass torts and the punitive damages that courts might see as legitimate (Underwood, 2009) are coming to weigh pretty heavily on groups seeking compensation or restrictions on actions that lead to dignitary, or reputational harms.

There are additional problems that relate to the use of torts in an effort to control the deployment of analytics in cyberspace.

6. Is It a Product, or Service, or Neither?

Part of the challenges we would face in relying on the tort system involves a determination of just what kind of “thing” this discriminatory technology actually is. It matters whether we are talking about a product, or a service, or something else entirely (Scott, 2008).

Product liability claims require three components: a product, a defect, and some harm (Geistfeld, 2009). All of these will be difficult for plaintiffs to identify convincingly with regard to the kinds of discriminatory technologies that are the focus of our concerns.

6.1. Software

First of all, we are talking about applications software. On the one hand, many of its features, including the ease with which it can be deployed simultaneously in an almost unlimited number of devices and locations makes software’s potential for causing harm so massive.

At the same time, part of its special character is that software has come to be routinely deployed even though its designers know that there are problems, or bugs that have not been corrected, or even fully identified (Scott, 2008). Over time we have come to expect, and perhaps even to accept these flaws and failures as an unavoidable fact of life, rather than a design defect for which someone should be held responsible (Callaghan & O’Sullivan, 2005; Childers, 2008; Rooksby, 2009).

6.2. Speech

Some of the difficulties we might face in our efforts to seek compensation for the harms caused by software applications is the fact that the technology and its outputs may actually be protected under the legal umbrellas that protect things that we define as speech.

For example, ratings agencies, from those that rate investments to those that estimate the size and composition of media audiences, have all managed to argue that their ratings are essentially opinions that have protection against state action under the First Amendment (Napoli, 2009). Although the ratings relied upon to justify the bizarre combination of mortgages into tradable securities have been recognized as wildly inaccurate, and perhaps fraudulent, American courts have not been willing to punish that form of speech (Freeman, 2009).

6.3. Tests

So, let’s turn to the question of tests. Tests are a way of life. We need to know how well we are doing in relation to others, or in relation to some standard of performance (Johnson, 2007). Generally, these tests are prospective in that they are used to predict how some system, or entity is likely to perform in the future. The consequences of error in these assessments are substantial.

For example, the rising cost of health care delivery reflects the increasing burden of providing just compensation for the hundreds of thousands of deaths, and millions of injuries that occur at the hands of medical personnel. Many of these injuries are the product of errors that are primarily informational. Many of these begin with an inaccurate diagnosis (Schiff et al., 2009). However, a good many of these involve errors in the interpretation of test results, rather than in a flaw in the design of the test.

Nevertheless, we want to be clear that tests will always have an error rate (Kane, 2006). They will also differ in terms of the kinds of errors that are consistently being made (Tatsioni et al., 2005; McGuire & Burke, 2008). We often talk about these in terms of Type I and Type II errors: finding problem or threat that is not actually there, or missing an indicator of risk that's as "clear as the bump on your head."

Although less attention is generally paid to the distribution of these errors, the fact is that some tests are biased in terms of their tendency to invite more Type I or Type II errors as a function of the particular attributes of the populations being assessed. Such errors may also reflect the influence of unwarranted (and often unacknowledged) assumptions about the circumstances governing the behavior of segments of those populations (Gandy, 2009, pp. 60-65).

Because tests are always going to exhibit some tendency toward error, we might have to define them as inherently dangerous products (Allhoff, 2009; Stewart, 2009).

According to the law in this area, inherently dangerous products are defined as those in which the danger "cannot be eliminated without impairing their intended function." It is in the very nature of discriminatory technologies that somebody is going to be harmed (Petty, 1998). The question that remains is about who will ultimately be blamed.

6.4. Mapping the Causal Path and Apportioning Responsibility

The manufacturer, the user, the subject of the use, and even so-called "innocent bystanders," who, perhaps were in the wrong place at the wrong time all have to be considered as potential sources of liability (Twerski & Grady, 2009; Henderson, 2009). Apparently each of them could be "apportioned" some share of the blame because of their contribution to an identified harm that occurs (Bubick, 2009; Van Velthoven & Wijck, 2009).

The list of potential parties in a civil suit doesn't end here. European policymakers are actively considering where to place responsibility for the harms that occur as the result of actions taken by autonomous digital agents or robots (Hildebrandt, 2008). It will be extremely difficult for courts to apportion responsibility when there are multiple computers, or programs, or agents involved in managing complex environmental systems across distributed networks (FIDIS, 2009b). These continually updated profiles and assessments of individuals' status will not be the product of any single actor or agent's behavior.

While there is certainly some responsibility to be assigned to the developer of the analytic software or algorithm, there is also responsibility for ensuring the quality of the data that are used in any particular analysis (Pottow, 2007). Of course, we might want to argue that each of the users of a technology have some identifiable "duty of care" with regard to the quality of the data they use in making important decisions about the kinds of opportunities that will be offered, or denied to some individual.

Apparently under tort law, there is also a requirement of "foreseeability" (Ausness, 2009; Twerski & Henderson, 2009; Epstein, 2010). There can be no doubt that most users of these technologies, and virtually all of the developers are aware of the fact that errors are almost certain when one relies on estimates based on the analysis of groups to guide decisions about any individual thought to be a member of a particular group.

Problems involved in the apportionment of responsibility are not insurmountable. However, addressing them tends to use up time and resources that might be better invested in figuring out how to minimize the harms in the first place.

6.4.1. Negligence vs. Strict Liability

Within the body of tort law in the US, there are also important distinctions that are drawn with regard to the mental state of a particular actor. Its assessment is part of a determination of whether the entity or actor has merely been negligent, or whether the harm was intentional.

Determining intent is a fool's mission. Intent can only be inferred, never proven, because there are no truly reliable sources for the facts that would reveal the truth. Instead, we may want to im-

pose “strict liability” on individuals and organizations that make use of software to guide discriminatory decisions. Users either know, or should have known that these systems are certain to generate a measurable, and perhaps even a predictable amount of harm. The imposition of liability is especially important if that harm tends to be concentrated upon particularly vulnerable segments of the population.

6.5. Assigning Value and Seeking Balance

Let us turn now to the nature of those harms, the specification of their value, and consideration of a basis upon which benefits and costs might have to be balanced against each other.

Much has been written about the nature of the harms we associate with invasions of privacy, or any of the recognized privacy torts (MacCarthy, 2010). There is a general sense that the harms are insubstantial, or minimal at best for the average person (Calo, 2010). In recognition of these limits, some regulations, such as the Video Privacy Protection Act have established “liquidated damages” that while minimal, mount up when a corporate actor harms a large number of persons (Solove & Rotenberg, 2003, pp. 554-558; Ludington, 2007).

Considerable attention has been focused on the harms associated with the errors that will always be made when decisions are taken on the basis of some test, algorithm or risk management scheme (Gandy, 2009). These are the harms that occur as the result of a misclassification, or biased assessments of value or worth.

Consider a fairly recent case in which 100 families sued the maker of a “Baby Gender Mentor” that was supposed to provide prospective parents with early identification of the sex of their baby to be. Allegedly, they failed to do so in these particular cases. The promoters of the service had claimed an incredible 99.9% level of accuracy. Unfortunately, the accuracy of a test does not ensure the accurate transmission of those results to the consumer (Novick, 2008).

In these cases, and in others like them, I wonder how courts should proceed to establish the value of the “losses” that occur when a prospective parent terminates a pregnancy on the basis of an inaccurate test, or the miscommunication of the results. What I am asking us to think about here is how we should proceed to estimate the lifetime value to the individual, to the family, and to society of a life forgone by mistake?

The American Law Institute’s latest attempt to get tort law right established a position of prominence for the need to balance costs or harms against the benefits that come with the use of some product, service, or technique (Twerski & Henderson, 2009). Naturally, there are no hard and fast rules about how these balances are supposed to be set (Gelman, 2009).

Since these decisions are increasingly being made in the context of prospective assessments of risk, we are going have to struggle over justifying an emphasis on low probability, but high consequence threats, versus those high probability, but low consequence risks that still tend to cumulate in troublesome ways (Allhoff, 2009; Lempert, 2009).

While dignitary harms are important for us to consider, they pale in comparison with the long term economic damages suffered by individuals, their families, and their communities. These damages arise when individuals have been refused opportunity, or have been victimized by predators who have been armed with information about just which people should be targeted with which particular sets of appeals (Pottow, 2007).

At the end of the day we have to conclude that the challenge of estimating the total value of the economic and non-economic harms suffered over time as a result of the cumulation of negative and erroneous assessments is massive. The fact that plaintiffs and their attorneys have to identify the multiple parties that bear joint and several liability for this accumulation of harm seems to me to be a burden that might be too heavy for society to bear.

Thus, we turn at last to consideration of the benefits of a regulatory response.

7. The Regulatory Response

We regulate markets, and market behavior, out of concerns for equity, as well as out of concern for efficiency (Ayres, 2007b). The fact that the impacts of design flaws are inequitably distributed should be at least one basis for justifying regulatory intervention.

Some economists invite us to think about these market failures in terms of externalities, the unintended byproducts, or consequences that fall upon third parties, rather than those directly involved in the production or consumption of particular goods and services. Ed Baker (2002, pp. 41-62) has identified several of the kinds of externalities that are associated with the way we organize various markets for information. Among the most important externalities that Baker identifies are those that influence audience member's interactions with other people (2002, pp. 46-47). If we extend this to include organizations and institutions involved in the provision of governance, finance, security and social services, the implications become clear.

We readily understand pollution as an externality, and we seem to agree that markets will not work on their own to insure the maintenance of healthy and sustainable environments. Thus, many agree that the regulation of pollution and other threats to the environment should be treated as explicit and important public policy goals. The same conclusion should be arrived at with regard to the information environment, wherein the protection of privacy would be just one of the specialized areas of concern (Strahilevitz, 2008).

The regulatory challenge then, would be to find ways to internalize the many external costs generated by the rapidly expanding use of analytics. That is, we need to find ways to force the providers and users of discriminatory technologies to pay the full social costs of their use.

7.1. Technology Assessment and Review

Success in this effort would mean that governments would have to engage more seriously in technology assessment. Just as privacy is a state of being that can be operationally defined in a variety of ways, the state of equality within a society can also be measured and evaluated against a standard, a goal, or an earlier moment in time (Bailey, 2008).

7.2. Aspects of Regulatory Control

The kinds of regulatory control systems that would have to be put in place will differ in terms of where authority, agency, and responsibility would be located. In one approach, we might assume that the regulators understand the nature of the problem, including the many complexities that exist along the causal chain (Mulherin, 2007). The regulator under this model would also have to know how to correct, or at least how to mitigate the harms associated with a given technology and the uses to which it is being put (Gerard & Lave, 2007).

An alternative model, one that some supporters characterize as being "performance based," leaves it up to the providers and users of technology to decide how they will use them. But then, of course, they would face penalties or taxes on the basis of the impact of their uses on the measured outcomes of interest (Sugarman & Sandman, 2008; Sugarman, 2009).

Think about how responsibility might be shared between producers and consumers of prepared foods in the context of an effort to manage the problem of obesity and its associated health effects (McCormick & Stone, 2007). The producers of the food, the marketers, the vendors, as well as the dietitians who set the menus for school lunches bear some responsibility. The parents and children do as well.

The problems involved in the apportionment of responsibility that emerge in the pursuit of tort claims, would also emerge as problems in the setting, imposition, and collection of fines or taxes. It seems there is no escape from the challenge of establishing a monetary value that would be required to compensate victims for their harms as well as any additional amounts that might be required to effect a behavioral change in those responsible for generating the harms in the first place.

Whatever regulatory model we choose, there must be a way to establish targets, or performance standards that can be used to adjust the administrative response to harms.

7.2.1. Safety and Performance Standards

The U.S. Consumer Products Safety Commission (CPSC) is concerned with protecting consumers, and their children from the risks of serious injury associated with the use of consumer products. Unfortunately, its requirements for product testing are far less demanding than those that have been established by the Food and Drug Administration (Polinsky & Shavell, 2010).

The Environmental Protection Agency (EPA) represents a special case of sectoral regulation in that its mission is framed in terms of the protection of health. However, because health is dependent upon the status of the environment in which we make our lives, the protection of the environment is central within its regulatory scope (EPA, 2006). Although the EPA is involved in establishing standards of performance for entities with a significant impact on the environment, it relies upon other agencies at the state level to enforce regulations that are assumed to more accurately reflect local conditions.

It is worth noting that environmental regulations in the US under the Clinton administration began to pay more attention to the distribution of the costs and benefits of particular technologies among disadvantaged populations. While the discourse of environmental racism was eventually translated into concerns about “environmental justice,” the underlying focus remained on equity and fairness in this domain (Taylor, 2000).

7.2.2. Regulating Discriminatory Technologies

Establishing regulatory limits on the use of discriminatory techniques has been extremely difficult in the United States, and progress in this area is likely to lag behind efforts within The European Community. Even in the face of substantial evidence that tests, measures and procedures have disparate impacts on groups that are already burdened by the effects of past discrimination, users of discriminatory technologies in the US have generally been able to avoid liability or restraint by merely claiming business necessity. Then they only have to offer token evidence of the appropriateness and comparative effectiveness of the systems being used (Ayres, 2007b).

Establishing enforceable standards for tests and measures with regard to error rates will also not be easy, especially with regard to dynamic fields in which new products and services are being developed and introduced on a continuing basis (Moses, 2007).

There has been some success in some narrowly defined markets in identifying kinds of data that cannot be used in the determination of eligibility, or in the setting of rates and prices (Doyle, 2007; Strahilevitz, 2008). In the case of race, gender and ethnicity, these forbidden measures have been identified at the national level as they apply in particular domains.

The recent establishment of limitations on the use of genetic information in determining access to health insurance and employment represents a pretty substantial leap forward toward limiting the use of information previously valued for its contribution to the estimation of future cost and benefit streams (Rothstein, 2008). We need to keep moving forward along this path toward placing restrictions on the use of information for the purpose of discrimination.

The regulatory approach being proposed here does not seek to ban the use of all discriminatory technologies. Indeed, we need to support the development and implementation of techniques that would help us to achieve greater equality more efficiently and effectively than we have been able to so far. In part, this means that the techniques that have been used primarily to minimize risks and avoid threats to profitability might be used in ways that maximize the distribution of opportunities to those who have been disadvantaged in the past.

For example, using tests to identify students from disadvantaged populations with above average levels of resilience, or other critical social skills would help to increase the number of such students who would benefit from admission to more challenging academic programs. Using these technologies to increase access to a range of opportunities for members of disadvantaged popula-

tions who have sought access would make substantial and beneficial improvements in levels of inequality without reducing levels of privacy and its correlates. It would remain to be discovered if improving levels of equality reduces profitability beyond levels that both economic theory, and ethical values would see as reasonable.

7.3. Specification of Requirements and Duties

Again, whatever regulatory model we choose, there has to be some expectation on the part of users of these technologies that there will be substantial costs associated with their failure to comply with the rules (Hoofnagle, 2010). Requirements to warn, or otherwise inform users and their customers about the risks associated with the use of these systems should not absolve system producers of their own responsibility for reducing or mitigating the harms. This is part of a common problem of determining where to place economic burdens or incentives as tools to shape behavior most efficiently and effectively.

Despite the problems that have emerged over determining the best way to implement a carbon tax as a way to reduce greenhouse gas emissions (GHG), the consensus suggests that we should impose the tax at the source – at the mines, the wells, and the refineries (Metcalf, 2009). This suggests that we should also place the initial burdens on the developers of these analytical systems, and leave it up to them to pass on the costs to their customers.

7.4. Policy Formation and Implementation

Of course, it is important for us to ask whether there is any compelling evidence that the regulatory path to the protection of privacy and its entailments is actually worth taking (FIDIS, 2009a).

This is an especially important question to ask in the context of what we have learned about regulatory capture and other strategic acts that have been used to limit the effectiveness of regulatory bodies (Freudenberg, 2005; Hoofnagle, 2010).

The difficulty in assessing the effectiveness of regulatory interventions reflects the complexity of the systems being regulated, the difficulty of deriving appropriate measures of performance, and the costs of gathering the necessary data (Sugarman, 2009).

Just as we have developed an array of specialized agencies and administrations concerned with protecting health and safety in relation to products, practices, and environments, we need to explore the creation of regulatory agencies with the responsibility for ensuring the well being, survival, and sustainable development of privacy in all of its forms. This regulatory goal should be pursued vigorously with regard to privacy's role in the reduction of social, economic, and political inequality within our nations (Wilkinson & Pickett, 2010).

8. The Political Challenge

We have reviewed our options with regard to pursuing our interests in limiting the use of discriminatory technologies within the regulatory regime described by Bennett (2011) and others. We concluded that considerable difficulties arise as a result of the tendency of this regime to emphasize the protection of data about individuals while ignoring the threats inherent in the use of data mining techniques to enable discrimination against members of statistically defined groups. We also concluded that consumers are largely incapable of exercising meaningful choice with regard to the rapidly expanding array of points within the matrix of networked interactions in which their interests will be placed at risk.

We also reviewed the options that might be available to individuals as members of groups seeking protection from, or compensation for the harms generated by the use of discriminatory technologies through the court system and the continually evolving law of torts. The limitations on the utility of torts are substantial, beginning with the special character of ratings and assessments as protected speech. Adjustments in the law that are directed toward establishing an appropriate bal-

ancing of interests between corporate actors on the one hand, and consumers on the other seem increasingly to favor one group over the other.

While acknowledging many of the limitations that exist within the regimes responsible for the regulation of technology, I still conclude that the regulatory path holds the greatest potential for protecting privacy as an aid to the enhancement of equality as a social goal.

Of course, the politics involved in developing this kind of regulatory response will be incredibly difficult and complex, in part because the numbers and variety of interests or “stakeholders” that will be affected is quite large (FIDIS, 2009a). Such a regulatory initiative risks being overwhelmed by appeals, claims, and demands that special exceptions be granted for industries, sectors, and initiatives that have already achieved some special status (Kinchy, Kleinman, & Autry, 2008; Hoofnagle, 2010).

Unfortunately, it is no longer clear what role public opinion can be expected to play in moving such an initiative forward. There is little value in framing arguments within this debate in terms of the usual unsympathetic victims (Gandy, 2009, pp. 191-194), because critically important audiences tend to quickly avert their gaze. Arguments have to be framed in terms of their importance to the larger collective.

This is not an impossible dream. The successful move against genetic discrimination in the US provides some lessons that we should not ignore. Genetic discrimination places us all at risk, and it seems likely that this fact explains its broad appeal as a policy target, at least in the United States. There may also be similarly broad appeals that can be derived from global efforts to protect the environment through the regulation of pollution and resource depletion.

While it is obvious that the EPA (EPA, 2006) is far from having achieved the goals that justified its creation, progress has been made, and it continues to be made against even more serious and immediate threats. Those of us seeking to negotiate sustainable futures for life on the earth tend to frame our appeals at a global scale, even while calling attention to the importance of local actions to reduce GHG. I believe that this emerging regulatory framework can be expanded to incorporate many of the concerns that we have identified with regard to social impact of discriminatory technologies

References

- Agre, P. (1997). Beyond the mirror world: Privacy and the representational practices of computing. In P. Agre & M. Rotenberg (Eds.), *Technology and privacy: The new landscape* (pp. 29-61). Cambridge, MA: The MIT Press:
- Allhoff, F. (2009). Risk, precaution, and emerging technologies. *Studies in Ethics, Law, and Technology*, 3(2), Article 2.
- Anderson, T. W., & Sclove S. (1986). *The statistical analysis of data*. Palo Alto, CA: The Scientific Press.
- Andrejevic, M. (2007). *iSpy: Surveillance and power in the interactive era*. Lawrence, KN: University Press of Kansas.
- Ausness, R. (2009). Product liability's parallel universe. Fault-based liability theories and modern products liability law. *Brooklyn Law Review*, 74, 635-667.
- Ayres, I. (2007a). *Super crunchers: Why thinking-by-numbers is the new way to be smart*. New York, NY: Random House.
- Ayres, I. (2007b). Market power and inequality: A competitive conduct standard for assessing when disparate impacts are unjustified. *California Law Review*, 95, 669-719.
- Bailey, J. (2008). Towards an equality-enhancing conception of privacy. *The Dalhousie Law Journal*, 31, 267-309.
- Baker, C. E. (2002). *Media, markets, and democracy*. New York, NY: Cambridge University Press.
- Bennett, C. (2011). In defense of privacy: the concept and the regime. *Surveillance & Society*, 8(4), 485-496.
- Callaghan, D., & O'Sullivan, C. (2005). Who should bear the cost of software bugs? *Computer Law & Security Report*, 21, 56-60.
- Calo, M. R. (2011). The boundaries of privacy harm. *Indian Law Journal*, 86(3), 1131-1162.
- Childers, S. (2008). Don't stop the music: No strict products liability for embedded software. *University of Florida Journal of Law & Public Policy*, 19, 125-184.
- Dahrendorf, R. (1979). *Life chances: Approaches to social and political Theory*. Chicago, IL: University of Chicago Press.
- Davenport, T., & Harris, J. (2007). *Competing on analytics: The new science of winning*. Boston, MA: Harvard Business School Press.
- Doyle, O. (2007). Direct discrimination, indirect discrimination and autonomy. *Oxford Journal of Legal Studies*, 27(3), 537-553.
- Epstein, R. (2010). Toward a general theory of tort law: Strict liability in context. *Journal of Tort Law*, 3(1), Article 6.

- FIDIS (2009a). Report D7.16 on *Profiling in financial institutions* [PDF document]. Retrieved from: <http://www.fidis.net/resources/deliverables/profiling/#c2597>.
- FIDIS (2009b). Report D17.3 on *Bridging the accountability gap: Rights for new entities in the information society?* [PDF document]. Retrieved from: <http://www.fidis.net/resources/deliverables/privacy-and-legal-social-content/>.
- Freudenberg, N. (2005). Public health advocacy to change corporate practices: Implications for health education practice and research. *Health, Education & Behavior*, 32(3), 298-319.
- Gandy, O. (2009). *Coming to terms with chance: Engaging rational discrimination and cumulative disadvantage*. Farnham, UK: Ashgate.
- Gandy, O. (2010). Engaging rational discrimination: Exploring reasons for placing regulatory constraints on decision support systems. *Ethics and Information Technology*, 12(1), 29-42.
- Geistfeld, M. (2009). The value of consumer choice in products liability. *Brooklyn Law Review*, 74(3), 781-805.
- Gellman, R. (2008). *Fair information practices: A basic history*. Retrieved from: <http://bobgellman.com/rg-docs/rg-FIPShistory.pdf>
- Gelman, L. (2009). Privacy, free speech, and "Blurry-Edged" social networks. *Boston College Law Review*, 50, 1315-1344.
- Gerard, D., & Lave, L. (2007). Experiments in technology forcing: Comparing regulatory processes of US automotive safety and emission regulations. *International Journal of Technology, Policy, and Management* 7(1), 1-14.
- Giddens, A. (1984). *The constitution of society*. Cambridge, UK: Polity Press.
- Gigerenzer, G., Swijtink, Z., Porter, T., Daston, L., Beatty, J., & Kruger, L. (1989). *The empire of chance: How probability changed science and everyday life*. New York, NY: Cambridge University Press.
- Gomez, J., Pinnick, T., & Soltani, A. (2009). Report *KnowPrivacy*. Berkeley, CA: UC Berkeley School of Information. Retrieved from: http://knowprivacy.org/report/KnowPrivacy_Final_Report.pdf.
- Grady, M. (2009). Unavoidable accident. *Review of Law and Economics*, 5(1), 177-231.
- Hand, D. (2006). Classifier technology and the illusion of progress. *Statistical Science*, 21(1), 1-15.
- Hayes, R. (1993). Measurement of information and communication: A set of definitions. In J. Schement & B. Rubin (Eds.), *Between communication and information* (pp. 81-103). New Brunswick, NJ: Transaction Publishers.
- Hildebrandt, M. (2008). Ambient intelligence, criminal liability, and democracy. *Criminal Law and Philosophy*, 2, 163-180.
- Hobart, M., & Schiffman, Z. (1998). *Information ages: Literacy, numeracy, and the computer revolution*. Baltimore, MD: The Johns Hopkins University Press.
- Hoofnagle, C. (2010). New challenges to data protection, Country Study B.1-United States of America. Part of the *Comparative study on different approaches to new privacy challenges in particular in the light of technology developments edited by D. Korff*. Brussels: European Commission, Directorate-General Justice, Freedom and Security. Retrieved from: http://ec.europa.eu/justice/policies/privacy/docs/studies/new_privacy_challenges/final_report_country_report_B1_usa.pdf.
- Johnson, V. (2007). Standardized tests, erroneous scores, and tort liability. *Rutgers Law Journal*, 38(3), 655-718.
- Kane, G. (2006). Field sobriety tests: Percentages of percentages, why "validation" studies fail to validate. *Trial Talk*, 5, 31-36.
- Kinchy, A., Kleinman, D., & Autry, R. (2008). Against free markets, against science? Regulating the socio-economic effects of biotechnology. *Rural Sociology*, 73(2), 147-179.
- Kritzer, H. (2007). Defending torts: What should we know? *Journal of Tort Law*, 1(3), Article 3.
- Lempert, R. (2009). *Low probability/high consequence events: Dilemmas of damage compensation*. Ann Arbor, MI: University of Michigan Law School Law & Economics Working Paper Series.
- Ludington, S. (2007). *Reigning in the data traders: A tort for the misuse of personal information*. Durham, NC: Duke Law School Working Paper Series.
- MacCarthy, M. (2010). *New directions in privacy: Disclosure, unfairness and externalities*. Privacy Law Scholars Conference at Georgetown University. Washington, DC: bepress.
- McCormick, B., & Stone, I. (2007). Economic costs of obesity and the case for government intervention. *Obesity Reviews*, 8(Suppl. 1), 161-4.
- McGeveran, W. (2009). Disclosure, endorsement, and identity in social marketing. *University of Illinois Law Review*, 2009(4), 1105-1166.
- McGuire, A., & Burke, W. (2008). An unwelcome side effect of direct-to-consumer personal genome testing. Raiding the medical commons. *JAMA*, 300(22), 2669-2671.
- Mello, M., Studdert, D., Thomas, E. J., Yoon, C. S. & Brennan, T. A. (2007). Who pays for medical errors? An analysis of adverse event costs, the medical liability system, and incentives for patient safety improvement. *Journal of Empirical Legal Studies*, 4(4), 835-860.
- Metcalf, G. (2009). Designing a carbon tax to reduce U.S. greenhouse gas emissions. *Review of Environmental Economics and Policy*, 3(1), 63-83.

- Mulherin, J. H. (2007). Measuring the costs and benefits of regulation: Conceptual issues in securities markets. *Journal of Corporate Finance*, 13, 421-437.
- Napoli, P. (2009). Audience measurement, the diversity principle, and the First Amendment right to construct the audience. *St. John's Journal of Legal Commentary*, 24(2), 359-385.
- Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford, CA: Stanford University Press.
- Nockleby, J. (2007). How to manufacture of crisis: Evaluating empirical claims behind "tort reform". *Oregon Law Review*, 86, 533-596.
- Novick, R. (2008). One step at a time: Ethical barriers to home genetic testing and why the U.S. health system is not ready. *Legislation and Public Policy*, 11, 621-649.
- Office of the Chief Financial Officer, EPA. (2006). *2006-2011 EPA strategic plan: Charting our course*. Washington, DC: United States Environmental Protection Agency.
- Ohm, P. (2010). Broken promises of privacy: Responding to the surprising failure of anonymization. *UCLA Law Review*, 57, 1701-1777.
- Petty, R. (1998). Interactive marketing and the law: The future rise of unfairness. *Journal of Interactive Marketing*, 12(3), 21-31.
- Phillips, D. (2004). Privacy policy and PETs. *New Media & Society*, 6(6), 691-706.
- Polinsky, A. M., & Shavell, S. (2010). The uneasy case for product liability. *Harvard Law Review*, 123, 1437-1492.
- Pottow, J. (2007). Private liability for reckless consumer lending. *University of Illinois Law Review*, 2007(1), 405-465.
- Rooksby, E. (2009). How to be a responsible slave: Managing the use of expert information systems. *Ethics and Information Technology*, 11, 81-90.
- Rothstein, M. (2008). Is GINA worth the wait? *Journal of Law, Medicine & Ethics*, 36(2), 174-178.
- Rubin, P., & Shepherd, J. (2008). The demographics of tort reform. *Review of Law and Economics*, 4(2), 591-620.
- Sandefur, R. (2008). Access to civil justice and race, class, and gender inequality. *Annual Review of Sociology*, 34, 339-58.
- Schiff, G., Hasan, O., Kim, S., Abrams, R., Cosby, K., Lambert, B. L., ... McNutt, R. A. (2009). Diagnostic error in medicine. *Archives of Internal Medicine*, 169(20), 1881-7.
- Scott, M. (2008). Tort liability for vendors of insecure software: Has the time finally come? *Maryland Law Review*, 67, 425-84.
- Solove, D. (2008). Data mining and the security-liberty debate. *The University of Chicago Law Review*, 75, 343-362.
- Solove, D., & Rotenberg, M. (2003). *Information privacy law*. New York, NY: Aspen Publishers.
- FTC (2010). *Protecting consumer privacy in an era of rapid change: Preliminary FTC staff report*. Washington, DC: Federal Trade Commission.
- Stewart, L. (2009). Strict liability for defective product design: The quest for a well-ordered regime. *Brooklyn Law Review*, 74(3), 1039-1059.
- Strahilevitz, L. (2008). Reputation nation: Law in an era of ubiquitous personal information. *Northwestern University Law Review*, 102(4), 1667-1738.
- Sugarman, S. (2009). Performance-based regulation: Enterprise responsibility for reducing death, injury, and disease caused by consumer products. *Journal of Health Politics, Policy, and Law*, 34(6), 1035-1077.
- Sugarman, S., & Sandman, N. (2008). Using performance-based regulation to reduce childhood obesity. *Australia and New Zealand Health Policy*, 5(26).
- Tatsioni, A., Zarin, D. A., Aronson, N., Samson, D. J., Flamm, C. R., Schmid, C., & Lau, J. (2005). Challenges in systematic reviews of diagnostic technologies. *Annals of Internal Medicine*, 142, 1048-1055.
- Taylor, D. (2000). The rise of the environmental justice paradigm. *American Behavioral Scientist*, 43(4), 508-580.
- Tenner, E. (1997). *Why things bite back: Technology and the revenge of unintended consequences*. New York, NY: Vintage Books.
- Tenner, E. (2004). *Our own devices: How technology remakes humanity*. New York, NY: Vintage Books.
- Tokson, M. (2011). Automation and the Fourth Amendment. *Iowa Law Review*, 96, 581-647.
- Twerski, A., & Henderson, J. (2009). Manufacturers' liability for defective product designs: The triumph of risk-utility. *Brooklyn Law Review*, 74(3), 1061-1108.
- Underwood, J. (2009). Road to nowhere or jurisprudential u-turn? The intersection of punitive damage class actions and the due process clause. *Washington and Lee Law Review*, 66, 673-807.
- Velthoven, B. V., & Wijck, P. V. (2009). Additive and non-additive risk factors in multiple causation. *Review of Law and Economics*, 5(1), 517-539.
- Vidmar, N., & Wolfe, M. (2009). Punitive damages. *Annual Review of Law and Social Science*, 5, 179-99.
- Wilkinson, L. (2008). The future of statistical computing. *Technometrics*, 50(4), 418-447.
- Wilkinson, R., & Pickett, K. (2010). *The spirit level: Why equality is better for everyone*. London, UK: Penguin.

- Wright, D., Gutwirth, S., Friedewald, M., Hert, P. D., Langheinrich, M., & Moscibroda, A. (2009). Privacy, trust, and policy-making: Challenges and responses. *Computer Law & Security Review*, 25(1), 69-83.
- Zimmer, M. (2008). The externalities of Search 2.0: The emerging privacy threats when the drive for the perfect search engine meet Web 2.0. *First Monday*, 13(3).

About the Author

Oscar H. Gandy, Jr.

is an emeritus professor of communication at the University of Pennsylvania (email: ogandy@asc.upenn.edu, website: <http://www.asc.upenn.edu/usr/ogandy/>). He is the author of: *The Panoptic Sort*, *Beyond Agenda Setting*, *Communication and Race*, and most recently, *Coming to Terms with Chance*. Earlier versions of this essay have been presented at the 2011 conference of the IAMCR, The Annenberg School for Communication, and the University of Ottawa Law School.