

IRRATIONAL PRIVACY?

JULIE E. COHEN*

I.	DO INFORMATION PRIVACY MARKETS WORK?.....	242
II.	WHAT WORK DO INFORMATION PRIVACY MARKETS DO?	245
III.	IS PRIVACY IRRATIONAL?	247
IV.	WHAT CAN BEHAVIORAL ECONOMICS TEACH US ABOUT INFORMATION PRIVACY?	249

The purpose of this symposium is to consider the implications of important recent work at the intersection of behavioral economics and information privacy. There is much to celebrate in this work, which has improved our understanding of the ways that markets in personal information actually function. Yet behavioral economic analysis has done little to defuse the most intractable problem that emerges from the study of so-called “information privacy markets.” The problem appears initially as a question of terminology. Discussing whether or not information privacy markets work presumes something rather important about what exactly such markets produce. The terminological question leads into an epistemological difficulty that has bedeviled information privacy law and policy: explaining what information privacy markets produce (if not “information privacy”) can cause information privacy claims to seem irrational at the most basic level.

This essay seeks to force the epistemological question onto the behavioral economists’ table. I will first argue that the term “information privacy markets” is a misnomer; information privacy markets do not in fact produce more information privacy but instead produce its opposite: more and more information about the actual or hypothesized attributes of individuals and groups. It does not follow, however, that a condition of information privacy (i.e., a condition of diminished or restricted access to personal information) is inherently antithetical to knowledge. Information privacy markets produce one kind of knowledge, but not the only kind. The condition of information privacy, and the kinds of knowledge it privileges, are best understood not as irrational but as differently rational—predicated on an approach to knowledge that we do not prize nearly as much as we should. In other words, information privacy’s seeming irrationality should trouble us a lot less than it seems

* Professor, Georgetown Law.

to. The problem, instead, is that we have not learned to name, value, and systematically pursue the kind of knowledge that a condition of information privacy produces. It remains to be seen whether behavioral economics can prove at all useful in exploring the sort of rationality that a commitment to information privacy represents.

I. DO INFORMATION PRIVACY MARKETS WORK?

We have come a long way from simplistic, uncritical models of information privacy markets—models positing that if people are provided with technological platforms from which to bargain over privacy preferences, they will do so without significant impediments. Over the decade or so, a number of talented and creative scholars have devoted their energies to exploring whether markets for consumer personal information work. For the most part, they have concluded that such markets are characterized by substantial failures. Some scholars and policymakers have seemed to expect that eventually, research on the behavioral economics of information privacy markets will uncover viable corrective strategies. That rather optimistic stance passes over some important questions about what a working market would look like and more fundamentally about whether “information privacy market” is a term that makes sense.

As currently constituted, markets for personal information have a number of large structural defects. At the front end of the process, where information is collected from individuals, most people have only a very limited understanding of the ways in which the information will be used.¹ In addition, many people will cheerfully disclose information about themselves to obtain particular transactional and relational advantages without pausing to consider the longer-term consequences.² At the back end of the process, where decisions must be made about the level of information security that proprietors of data reservoirs ought to provide, consumers lack both the information and the expertise to evaluate the results.³ In the world of Web 2.0 social networking platforms, people

1. See, e.g., Alessandro Acquisti & Jens Grossklags, *Privacy and Rationality in Individual Decision Making*, IEEE SECURITY & PRIVACY, Jan./Feb. 2005, at 26-33; Alessandro Acquisti & Ralph Gross, *Imagined Communities: Awareness, Information Sharing, and Privacy on the Facebook*, in PRIVACY ENHANCING TECHNOLOGIES: 6TH INTERNATIONAL WORKSHOP, PET 2006, CAMBRIDGE, UK, JUNE 2006, REVISED SELECTED PAPERS 36-58 (George Danezis & Philippe Golle eds., 2006); Alessandro Acquisti & Jens Grossklags, *What Can Behavioral Economics Teach Us about Privacy?*, in DIGITAL PRIVACY: THEORY, TECHNOLOGIES AND PRACTICES 363 (Alessandro Acquisti et al. eds., 2007).

2. See, e.g., Scott R. Peppet, *Unraveling Privacy: The Personal Prospectus and the Threat of a Full Disclosure Future*, 105 NW. U.L. REV. 1153, 1157-58 (2011) (citing Robert H. Frank, PASSIONS WITHIN REASON 104 (1988)).

3. Danielle Keats Citron, *Reservoirs of Danger: The Evolution of Public and Private Law at the Dawn of the Information Age*, 80 S. CAL. L. REV. 241, 267 n.142 (2007).

also eagerly participate in profiling others for a variety of purposes, some social and others punitive.⁴

Even more importantly, dysfunction in privacy markets has a dynamic aspect. Over time, people can be expected to (over)disclose more and more information, both because they have become inured to disclosure and because the equilibrium in the marketplace has tipped toward disclosure as a condition of market entry, a dynamic that Scott Peppet terms “unraveling.”⁵ (I cannot resist noting that this dynamic has long been thoroughly appreciated in the literature on surveillance and has made its way into the economic literature on privacy only recently, but I’m nonetheless heartened to see its arrival.)⁶ The incentives toward (over)collection and (over)processing are similarly robust. Lior Strahilevitz’s work at the intersection of privacy and antidiscrimination protection has suggested that, to borrow the old chestnut about Internet censorship, profiling routes around attempts to disrupt it.⁷

Each of these market dynamics can itself be subjected to more economic analysis, which is to say that we can develop good models of the level and direction of anticipated privacy market dysfunction over time. For example, one might imagine a game theoretical analysis of the effects of unraveling, or of efforts to use information to discriminate in employment. Such work might yield useful prescriptions about the need to require still more thorough disclosure or the kinds of prohibitions that an enlightened regulator might reasonably expect to be effective given the reality of self-interested marketplace behavior.

Behavioral economic analysis of markets in personal information is and will remain an important tool for legal scholars and policymakers. It enables us to get traction on important issues relating to the incentives and the bounded rationality of market participants. One might be entirely justified in thinking that, over time, it could help us address a variety of practical problems that concern how markets for consumer personal information might be made to work both more transparently (as to the participants) and more securely (as to malicious third parties). If there are going to be information privacy markets, and there surely are, it’s useful to have effective rules of the road.

But now we come to the nub of things. What are information

4. Lior J. Strahilevitz, *Reputation Nation: Law in an Era of Ubiquitous Personal Information*, 102 NW. U. L. REV. 1667, 1670 (2008).

5. Peppet, *supra* note 2, at 1176.

6. See, e.g., MARK ANDREJEVIC, *ISPY: SURVEILLANCE AND POWER IN THE INTERACTIVE ERA* 5-6 (2007); Kevin D. Haggerty & Richard V. Ericson, *The Surveillant Assemblage*, 51 BRITISH J. SOCIOL. 605, 613-14 (2000).

7. Lior J. Strahilevitz, *Privacy versus Antidiscrimination*, 75 U. CHI. L. REV. 363, 373-75 (2008); see Philip Elmer-Dewitt, *First Nation in Cyberspace*, TIME, Dec. 6, 1993, at 64 (“[A]s Internet pioneer John Gilmore puts it, ‘The Net interprets censorship as damage and routes around it.’”).

markets that “work” properly supposed to look like? Answers to this question remain incoherent on the most basic levels: We may think we know what is being bought and sold, but who are the buyers and who the sellers? And what do information privacy markets produce?

Identifying buyers and sellers is surprisingly tricky because information privacy markets have a business-facing side and a consumer-facing side. The ultimate source of the information is the individual consumer, but that does not make the consumer the customer as far as most market participants are concerned. The website for TRUSTe, a leader in the emerging market for providing cutting-edge privacy services, makes clear that the “customers” are the businesses that want to obtain certification services attesting to the reliability of their representations about their information processing practices.⁸ Here what is being bought and sold is a trusted brand. The websites operated by data aggregators such as LexisNexis Risk Solutions (which acquired ChoicePoint) tout a variety of business-to-business services and do not appear to contemplate business-to-consumer relationships at all.⁹ As applied to both of these examples the term “information privacy markets” introduces a certain amount of obfuscation about the primary locus of the market activity, and about how the casual observer might reliably identify the participants and interpret their economic and political interests.

Matters become even more curious when we turn to the question of production. In the ordinary understanding of the word, a “market” for something is an exchange network within which that something is produced and traded. Markets for apples produce and facilitate trade in apples; markets for widgets produce and facilitate trade in widgets, and so on. Applying this ordinary rule of interpretation might lead us to conclude that information privacy markets produce and facilitate trade in information privacy. If so, we could all breathe a sigh of relief and proceed to have the usual debates about whether and to what extent regulation will produce more privacy or better privacy.

Yet if one conceives of a condition of information privacy as a condition in which there remain important and durable gaps in the information about oneself that is accessible to others, the conclusion is inescapable that information privacy markets do not exist to produce more information privacy. They exist precisely so that over time we may have less information privacy, and the disagreements in the economic literature concern merely the details about how the condition of diminished information privacy will be structured. Within the framework

8. See TRUSTe, <http://www.truste.com> (last visited May 19, 2012).

9. See *Risk Solutions*, LEXISNEXIS, <http://www.lexisnexis.com/risk> (last visited May 19, 2012).

of consumer protection that has come to dominate discussions about markets in personal information, use of the term “information privacy markets” arguably is the ultimate deceptive act.

II. WHAT WORK DO INFORMATION PRIVACY MARKETS DO?

If information privacy markets don't produce information privacy, what do they produce? There is an important extent to which all of the important work about the behavioral economics of markets in personal information is a distraction from this most basic of questions. Here reliance on a methodology that is essentially technocratic—that emphasizes defects in market calibration rather than overarching policy objective—has its costs. Work in the technocratic market-calibration mode can be a way of avoiding more difficult questions about what all of the “sorting” and “signaling” and “unraveling” are good for.

At the most general level, sorting and signaling operationalize a system for understanding what consumers are willing to pay and, for more complex goods that entail ongoing allocation of risk, what terms providers are willing to offer. The purpose of sorting and signaling, in other words, is to differentiate—to discriminate, if one prefers—among groups of consumers according to frameworks that seem to make sense based on the various goals of market participants. These activities generate both costs and wealth, depending on how they are conducted and on the regulatory parameters put in place to channel and constrain them. Somewhat more precisely, sorting and signaling work in conjunction with heuristics designed to identify and systematize consumers as sources of both profit and risk. The heuristics are themselves informed by preconceived ideas about risk and about sensible ways of sorting. They are both empirically supported and socially constructed.

Scholarship in the technocratic market-calibration mode has made some useful contributions to our understanding of the heuristics behind sorting and signaling. In general, however, scholarly investigations into the ways that we sort and signal have proceeded narrowly along the fault lines generated by liberal anxieties about autonomy and paternalism. Those anxieties suggest certain types of objections to sorting and signaling, which I will call the diminished autonomy and invidious discrimination objections, respectively. We might not want to respect people's preferences to disclose certain kinds of information if the people in question lack the capacity for informed choice. Similarly, we might not want to respect the preferences of consumer-facing businesses to condition the provision of goods or services on particular disclosures, or to respect the preferences of employers to hire only certain kinds of people, if the preferences simply reflect bias and are harmful to society.

Both the diminished autonomy objection and the invidious discrimination objection to markets in personal information seem to dictate limited interventions aimed at narrow sets of particularly egregious practices. Such interventions operate externally to information privacy markets, acting as situation-specific correctives. Their scope is cabined by a background presumption of fully-fledged autonomy, which suggests that we ought to let people make their own choices most of the time. It is also cabined by assumptions linking increased empirical support with increased rationality. Those assumptions suggest that we ought to be able to determine which kinds of discrimination are invidious and which legitimate, and that ordinarily, markets fueled by information will operate to discipline the choices that businesses and employers make.

Now, however, consider two more fundamental objections that we as a society might have to processes of sorting and signaling, and in particular to the heuristics behind them: The first is what we might call the economic justice objection—i.e., that sorting and signaling work to operationalize a system characterized by preferential treatment for the wealthy and the maximum extraction of consumer surplus from everyone else. The second objection, which I will call the capabilities objection, has to do with breathing room for evolving subjectivity—with the claim that all human beings, whatever their resources, require a baseline degree of freedom from categorization in order to flourish as human beings.¹⁰ If these objections are valid, they signal a need for a comprehensive reevaluation of markets for personal information, and for a much more extensive menu of market interventions.

We are now standing at the brink of an epistemological wormhole. The surest indication that we are approaching an anomaly in the philosophical continuum is that it can be enormously difficult to recognize the economic justice objection and the capabilities objection as legitimate objections to personal information processing at all. Recall that the epistemological difficulty I identified at the start has to do with the presumed equivalence between the processing of personal information and the production of knowledge. Both the economic justice objection and the capabilities objection suggest that broad restrictions on the processing of personal information might be desirable. But isn't the idea of forms of prohibited knowledge about people contradictory to everything that post-Enlightenment Western culture has stood for? And isn't it exactly the point of markets to separate people from their disposable surplus in exactly the proportion that they have it? Here information privacy claims stand in seeming opposition to an entire way

10. JULIE E. COHEN, CONFIGURING THE NETWORKED SELF: LAW, CODE, AND THE PLAY OF EVERYDAY PRACTICE 148-50 (2012).

of thinking.

III. IS PRIVACY IRRATIONAL?

To pose these questions in the way I have just done is to suggest an answer that is sociological (or sociocultural) rather than philosophical: The activities of sorting and signaling do not simply enable a particular type of decision-making, but also enact a set of deeply held beliefs about the value of information and its relationship to knowledge and understanding. Those beliefs hold that particularized sorting is the route to the most accurate knowledge, the most frictionless and responsive markets, and the fairest distribution of benefits and penalties. This philosophy of knowledge is so deeply rooted as to be almost invisible, but it is a culturally situated philosophy nonetheless. To grapple with the problem of whether information privacy claims are as deeply irrational as they can sometimes appear, we must bring the almost-invisible into critical focus.

Once again it is useful to consider definitions: Rational actions and beliefs are “guided by reason, principles, fairness, [or] logic;” irrational decisions and beliefs are not.¹¹ Perhaps, a commitment to information privacy expresses irrationality, but there is at least one other possibility. A commitment to information privacy might simply reflect a different understanding of what reason, principles, fairness, and/or logic require.

What evidence do we have that the second possible conclusion might actually be the correct one? One kind of evidence concerns the transformative properties of information. We know that the meaning of information is not constant, but rather changes when it is aggregated, and again when it is shared. Where we have gone wrong, I think, is in seeing only the positive side of this. A philosophy of knowledge that prizes ever-increasing particularization denies the possibility of a negative information externality. We are primed to reject the idea that information externalities could be like pollution externalities—i.e., transformative of the environment in an undesirable way—and to think instead that the problem is the “short attention spans” that prevent us from properly weighing and valuing new data.¹² That diagnosis permits us to remain agnostic about the best policy solution; in theory, more data, more data processing capability, and better heuristics could cure the attention deficit that plagues mere humans. But if information externalities are like pollution externalities, curing the attention deficit will make matters

11. MERRIAM-WEBSTER'S DICTIONARY OF LAW 400-01 (2011); *see also* OXFORD AMERICAN DICTIONARY OF CURRENT ENGLISH 661 (1999).

12. *See* JEFFREY ROSEN, THE UNWANTED GAZE: THE DESTRUCTION OF PRIVACY IN AMERICA 200-01 (2000).

worse, not better.¹³

Another kind of evidence relates to the slipperiness of efforts to isolate and contain purely normative dimensions of information policy. Antidiscrimination policy is an example of the containment dynamic at work. We're comfortable concluding that a particular level of detail isn't relevant, normatively, to views that we have chosen to hold about how people should be treated, but the comfort lasts only so long as the prohibition that effectuates an antidiscrimination norm is understood to be an anomaly. In information policy, however, the normativity goes all the way down.

A good example of the normative embeddedness of information policy is the Genetic Information Nondiscrimination Act (GINA), which embodies a conclusion that it is somehow more wrong for employers and insurers to sort prospective employees and insureds based on genetic markers than it is to sort based on other kinds of medical information, even though conditions that result from genetic predispositions undeniably produce both private and social costs.¹⁴ GINA reflects a judgment about individual accountability that is normative as well as descriptive. Another example is the Fair Credit Reporting Act, which specifies elaborate procedures to be followed in the disclosure of consumer reports but has almost nothing to say about the sorts of information that may be considered or about the substantive obligations of a consumer reporting agency to treat individuals fairly.¹⁵ Here the presumption is that individuals are accountable for every variable that a credit provider might want to consider. That too is a judgment that combines normative and descriptive aspects.

If the activities of information collection and processing are never only descriptive but always partly normative—always laden with culturally-embedded values about accountability and its effects on risk allocation—then it makes much less sense to rail against the imposition

13. One might hypothesize that making personal-information entitlements tradable and more expensive (as privacy regulations of the notice-and-choice variety tend to do) could buy time for the invention of “clean” information processing technologies (as the argument for tradable pollution entitlements holds), but it is also possible that the “more information is better” ideology will operate to prevent formation of the appropriate incentives.

14. Genetic Information Nondiscrimination Act of 2008, Pub. L. No. 110-233, 122 Stat. 881 (2008).

15. See Fair Credit Reporting Act, 15 U.S.C. §1681b (establishing parameters governing disclosure of consumer reports); §1681c(a) (excluding from consumer reports certain types of information relating to prior judicial proceedings and tax liens); §1681e(b) (requiring consumer credit agencies to follow “reasonable procedures to assure maximum possible accuracy”); *Sarver v. Experian Info. Solutions*, 390 F.3d 969, 972-73 (7th Cir. 2004) (holding that statutory reasonableness requirement does not mandate separate examination of each computer-generated report for anomalous information); Elizabeth De Armond, *Frothy Chaos: Modern Data Warehousing and Old-Fashioned Defamation*, 41 VAL. U.L. REV. 1061, 1101-14 (2007).

of other value-based limits simply on the basis of their supposed irrationality. The conversation about information privacy might then become a conversation about what kinds of knowledge are possible and which ones we ought to value. Within such a conversation, limits on information processing might come to seem desirable and even essential to sound information policy.

IV. WHAT CAN BEHAVIORAL ECONOMICS TEACH US ABOUT INFORMATION PRIVACY?

If the account of personal information processing that I have offered is at least plausible, behavioral economic models of personal information markets confront a new challenge: that of determining how to model the information externality. I suspect that meeting this challenge will be extraordinarily difficult. Behavioral economics doesn't eliminate the presumption of the autonomous subject that is conventional within economic theory, but only relaxes ancillary assumptions about the subject's perfect rationality. Because the information externality operates directly on subjectivity, producing not just different (better informed) decisions but ultimately different subjectivity, a different approach to modeling may be needed. Whether behavioral economics, or any other sort of economics, can supply the appropriate tools remains to be seen.

Alternatively, it may simply be that we need to learn to make (some) public policy decisions about information privacy without economic modeling or even the possibility of economic modeling—in other words, to make non-data-driven decisions about what limits ought to be placed on our insatiable urge to make data-driven decisions. This could well be done in the service of the sort of rationality that I have described. The question is whether we have the wherewithal to do it.

