Self-surveillance Privacy

Jerry Kang*
Katie Shilton
Deborah Estrin
Jeff Burke
Mark Hansen

Version 1.2

January 28, 2011

© by authors

Please do not quote, cite, copy, or distribute further,
without explicit permission.

---

* Professor of Law, UCLA School of Law, Professor of Asian American Studies (by courtesy), Korea Times--Hankook Ilbo Chair in Korean American Studies. ‹kang@law.ucla.edu› ‹http://jerrykang.net›. All other authors affiliated with UCLA CENS (Center for Embedded Networked Sensing).

Research assistance provided by: Jonathan Feingold, Emily Reitz, and Isaac Silverman.

Presented at: Privacy Law Scholars Conference 2009, UCLA School of Law Summer Workshop Series 2009.

Supported in part by: UCLA Academic Senate, UCLA School of Law, Center for Embedded Network Sensing.

Helpful comments by: Richard Abel, Ann Carlson, Jeff Jonas, Sung Hui Kim, Ken Klee, Tim Malloy, Jon Michaels, Stephen Munzer, Helen Nissenbaum, Marc Rotenberg, Noah Zatz, and Eric Zolt.

Introduction 1
I. Self-Surveillance 2
A. New Technologies 2
B. Benefits 8
II. Threshold Clarifications 11
A. Privacy Metrics 11
B. Privacy Displacement 14
C. Distinguishing Harder Privacy Problems 17
1. Not Third-party Surveillance of Us 18
2. Not Our Surveillance of Third-parties 19
III. Personal Data Guardian 21
A. Personal Data Guardian 22
B. Personal Data Vault 23
C. Legal Relations 27
1. Fiduciary Duties 28
2. Evidentiary Privilege 29
IV. Objections 36
A. Implausible? 36
B. Useless? 40
1. Third-party Surveillance 40
2. Genie out of the Bottle 41
Conclusion 46

INTRODUCTION

[1] It has become cliché to observe that new information technologies endanger privacy. Typically, the threat is viewed as coming from Big Brother (the government) or Company Man (the firm). But for a nascent data practice we call "self-surveillance," the threat may actually come from ourselves. Using various existing and emerging technologies, such as GPS-enabled smartphones, we are beginning to measure ourselves in granular detail - how long we sleep, where we drive, what we breathe, what we eat, how we spend our time. And we are storing these data casually, perhaps promiscuously, somewhere in the "cloud," and giving third-parties broad access. This data practice of self-surveillance will decrease information privacy in troubling ways. To counter this trend, we recommend the creation of the Privacy Data Guardian, a new profession that manages Privacy Data Vaults, which are repositories for self-surveillance data.

[2] In Part I, we describe the emerging data practice of self-surveillance, which has been enabled by various new measurement and communication technologies. We explain how self-surveillance can produce substantial benefits to both the individual and society, in both intrinsic and instrumental terms. Unfortunately, such benefits may never be achieved without substantial privacy costs.

[3] Part II makes threshold clarifications about those privacy costs. It proffers two different metrics by which privacy might be measured and explains why the rise of self-surveillance will entail the net loss of privacy under either metric. We also point out that the problem of self-surveillance (our surveilling us) is, fortunately, more tractable than related privacy problems, such as third-party surveillance of us and our surveillance of third-parties.

[4] Having cleared this brush, we turn to our central proposal-the creation of the Personal Data Guardian, a professional whose job it is to maintain a client's self-surveillance data in a Personal Data Vault. In addition to providing technical specifications of this approach, we outline the specific legal relations, which include a fiduciary relationship, between client and Guardian. In addition, we recommend the creation of an evidentiary privilege, similar to a trade secret privilege, that protects self-surveillance data held by a licensed Guardian.

[5] Finally, Part IV answers objections that our solution is implausible or useless. We conclude by pointing out that various legal, technological, and self-regulatory attempts at safeguarding privacy from new digital, interconnected technologies have not been particularly successful. Before self-surveillance becomes a widespread practice, some new innovation is needed. In our view, that innovation is a new "species," the Personal Data Guardian, created through a fusion of law and technology and released into the current information ecosystem.

I. SELF-SURVEILLANCE

A. New Technologies

[6] Bloggers and web masters are familiar with Google Analytics- a widely-adopted set of visualization tools that support examination of website traffic patterns.[1] A script sends website visitor data to Google, which then analyzes the traffic patterns with remarkable granularity and provides results through flexible visuals. One can easily see the IP address of who has visited, from where (geographically and from which prior page), when, how often, for how long, and through which keyword search. It's also free of charge.

[7] What's interesting is that new technologies allow us to cull, then analyze, similar sorts of details about not only our websites but also ourselves. Here are three examples. *RescueTime.com* allows the installation of a tiny software application that tracks how we spend time on our computer, down to the second.[2] If you want to know how much time you waste surfing particular Web sites, on an average Monday, you can easily collect that data.

---

[1] *See* ‹http://www.google.com/analytics/› (last visited December 21, 2010).

[2] As of December 2010, Rescuetime advertises two products, Pulse and Empower. The Pulse product is for "employee tracking" by management; in this sense, it is old-school surveillance. By contrast, the Empower product is more for self-analytics in that an individual voluntarily initiates the data collection for self-analysis. But even in this context, the meaning of the data collected turns on "peer" comparisons. See ‹http://www.rescuetime.com/› (last visited December 21, 2010).

[8] GPS manufacturer Garmin's motionbased.com is a web application that records our location in order to analyze outdoor training and fitness regimens.[3] If you are curious how long your typical morning jogs are, and whether you are improving your pace, it is now simple to collect that information and analyze it.

[9] Finally, Fitbit is a tiny piece of hardware that can be clipped on your clothing, which measures how many steps you've taken, how active you have generally been, and how many calories you have burned. In addition, it can track your sleep, and all of these data are uploaded wirelessly to their web site, which provides pretty graphs of the day and night's activity level.[4] A similar device called DirectLife, from Philips, includes data analysis and coaching from fitness and nutrition experts.

[10] These examples portend the rise of "self-surveillance"-- a data practice that measures, collects, and stores self-surveillance data. Self-surveillance data, in turn, are measurements of the self, initiated by the self, for the primary purpose of measuring the self, using sensors that are in one's control. By "measurements of the self", we mean a recording (fixed expression) of an observation about the self, which may include the environment to which the self is exposed. These data include metadata about the data recorded, such as the time and place of the sensing moment. By "in one's control," we mean that these devices are under a person's direct physical control, such as a heart rate meter that stores data onto local flash memory. They could also be under more indirect control, the degree to which could be measured by the ease with which the person can simply

---

[3] *See* ‹motionbased.com›. See also Personal projects: Daytum, http://daytum.com/ (makes use of Google charting API); Mycrocosm, http://mycro.media.mit.edu/ "a web service that allows you to share snippets of information from the minutiae of daily life in the form of simple statistical graphics"; Me-trics (pulls data from rescuetime and twitter and others) http://beta.me-trics.com/ and looks for correlations (yes, a little statistics); moodstats http://www.moodstats.com/ ; Nike+ http://nikeplus.nike.com/nikeplus/index.jhtml; Nathan Yau's Your Flowingdata http://your.flowingdata.com/ (providing wide open flexibility over data that can be tracked).

[4] *See* David Pogue, *Getting Fit with 2 Bits of Help*, New York Times, December 16, 2009.

turn off data collection without large transaction costs or loss of services from third parties.[5]

[11] Self-surveillance data includes, but is not restricted to, data collected through non-subjective and automatic sensors. By "non-subjective", we mean that they record data, such as location or acceleration, without asking for subjective introspection or self-reports from the individual. Also, these are "automatic" in that they collect data in a set-it and forget-it mode, which after initial configuration by the individual does not require manual input of information on an incident-by-incident basis.

[12] Non-subjectivity and automaticity make it more likely that huge datastreams will be collected invisibly in the background. That said, these features are not strictly necessary in a definitional sense for self-surveillance. For instance, we would count as self-surveillance the Experience Sampling Method developed by Mihaly Csikszentmihalyi,[6] which roughly involves an individual carrying a device that prompts her for self-reports about her status, such as happiness--even though the answers

[5] "Self-surveillance" is an odd term, and "self-monitoring" could be used in the alternative. We use the more jarring phrase because "surveillance" evokes greater threat, which we think is warranted given the privacy stakes. Moreover, we want to question the psychological and philosophical assumption that a person is so unified and internally consistent, especially over time, that the idea of surveilling oneself seems silly, as if we had to keep an eye on our own left hand lest it do something bizarre or inappropriate. For example, at a single moment, a person may have conflicting desires - think about wanting dessert (when we don't want to want it) or avoiding exercise (when we want to want it). In such cases, it seems reasonable to suggest that one part of the self is surveilling another part, in order to constrain or facilitate certain behaviors. The point is more vivid if we think about moments separated in time. Imagine that Johnny "consented" at the age of 16 to disclose certain images or facts on the Internet. Now, at the age of 30, Johnny regrets those disclosures but can't delete the information from public view. Johnny (present) is, of course, considered to be the same person as Johnny (past). Yet it seems reasonable to suggest that Johnny (past) has bound Johnny (present) through certain information choices made previously. Put another way, Johnny (present) is subject to a sort of data surveillance inflicted upon him by Johnny (past).

[6] Mihaly Csikszentmihalyi, Flow ( ).

rely on self-reports and not external measurements. Similarly, we would count as self-surveillance a calorie counting practice that requires the individual to photograph manually all food eaten even though such data capture isn't automatic (in the sense that taking the picture requires manual actuation at meal events).[7]

[13] *Participatory Sensing example*. To make our discussion more concrete, we explore a specific case study of self-surveillance—Participatory Sensing developed at UCLA's Center for Embedded Network Sensing (CENS). Participatory Sensing (PS) coordinates mobile devices such as smartphones for use as self-surveillance, personal wellness, and research instruments. To support the widest audience of users, Participatory Sensing uses off-the-shelf mobile phones[8] running specialized software.[9] The software

[7] With any definition, there will be hard cases. For example, one could analyze one's eating patterns by examining one's credit card transactions. Should, then, using a credit card be considered "self-surveillance" and the transactions listed on a monthly credit card bill deemed "self-surveillance data"? We think this would be at the fringes of the definition. Most important, the credit card data are collected for the primary purpose of facilitating a credit transaction and accurate billing, not for measuring oneself. In addition, it's not clear that the credit card transaction network should be considered to be in one's direct or indirect control, given our definition.

[8] Widespread penetration and use of mobile phones makes them attractive tools for participatory sensing and other types of self-monitoring. These always on, always present devices can capture locations and context information, infer habits and routines, and provide detailed, individualized assessment of behavioral and environmental factors.

[9] Participatory sensing is inspired by, and draws its name from, the broader tradition of *participatory research* (PR). PR is a set of methods that position research subjects as co-investigators (Cargo & Mercer, 2008). PR traditions develop their research questions with the cooperation of partner communities and engage community members in research design, implementation, analysis, and dissemination. Involvement with every stage of the research process allows participants to target local knowledge and benefit from the results of systematic investigations. PR successes in health and environmental research have improved the ability of marginalized or underserved

collects data using phones' available onboard sensors: cameras, a microphone, GPS or cell tower location, accelerometers, user-prompted entry, and Bluetooth connections to other devices. The software then uploads the geo-coded and time-stamped data to a server that performs data processing, aggregation, and modeling, and displays the results to each user via private web interfaces.

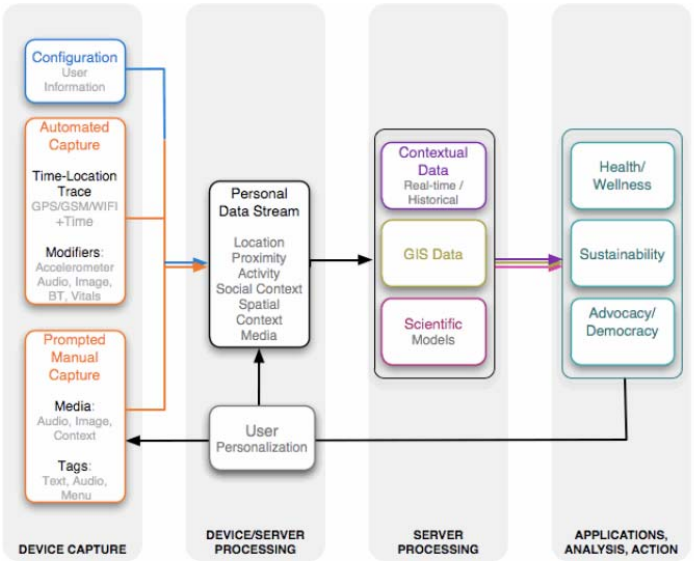

Figure 1. Participatory Sensing Processes

[14] Participatory Sensing relies on a series of processes, as shown on Figure 1. Smartphones automatically record the time, location, and activities of a participant by sampling GPS or cell tower location. The "Personal Data Stream" captured by the mobile device is then automatically uploaded to secure servers via the wireless mobile phone network. The server processes the data using models to estimate participant activities, for example using location and velocity to determine whether the individual is walking, running, biking, or driving.

[15] Traces that combine time, activity, and location can also support various health applications. For example, changes in

groups to act on the results of the data they have helped collect and analyze (Horowitz et al., 2009).

work, sleep, and weekend activities can serve as indicators of fatigue, depression, or increasing side effects. Similarly, features of location-activity traces, such as how much, how quickly, and how far a person walks outdoors, can serve as outcome markers for treatment of neuromuscular diseases or rehabilitation from stroke or surgery. Specific aspects of health status, such as pain, side effects, physiological self-measurements, and medication adherence patterns, can also be sampled using the Experience Sampling Method described above. For example, the smartphone can prompt the user to check and enter a physiological parameter (e.g. blood glucose), or a perception such as dizziness level. The mobile phone geo-codes and timestamps these responses and uploads them to the individual's data store to create an additional time series. The server can also link the data to web-based Geographic Information Systems (GIS) that, for example, document environmental hazards such as air pollution, name places and contexts (e.g., bars, home), and document characteristics of a community or social environment.

[16] Third party application service providers (3P-ASPs) can create the processing and models necessary to begin interpretation of Participatory Sensing data. For instance, CENS projects have included a wellness application that helps users discover when and where they engage in eating that's "off plan" or different from their objectives; a health application that helps chronic illness sufferers track relationships between medicine adherence, side effects, and personal mobility;[10] a project for biking commuters to collect and compare their cycling routes; and the Personal Environmental Impact Report (PEIR)[11]-- an application that gives users daily feedback on both their carbon footprint and their exposure to air pollution.[12] Participants can

[10] http://andwellness.cens.ucla.edu

[11] http://peir.cens.ucla.edu/

[12] Although we have highlighted self-surveillance examples, other CENS PS projects extend far beyond this scope. Indeed, the PS research was launched initially to support participatory sensing activities in which people decide what, how, and when to sense-not only themselves but features of the world around them, to collect and analyze geo-tagged imagery in support of ecological, public health, and cultural goals. For example, data can be collected as part of an explicit campaign, undertaken by many users in collaboration. One such

use these models as jumping off points for their own interpretations, or in even more participatory projects, collaborate with application providers on new model creation.

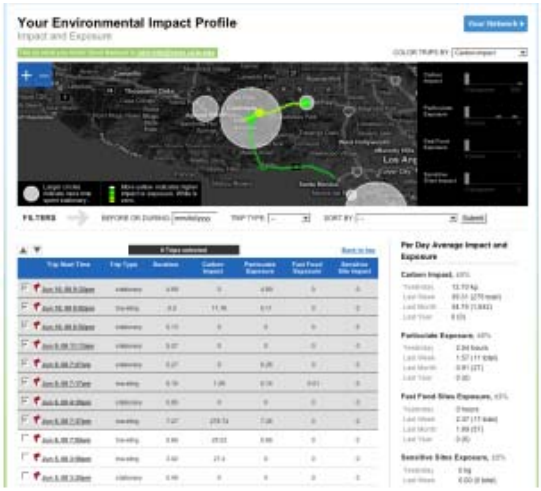

Figure 2. Environmental Impact

B. Benefits

[17] People engage in self-surveillance out of a natural interest in their selves. The personal data collected can then be processed to produce self-knowledge[13] that has instrumental and intrinsic value, not only to individuals but also to society.

[18] *To individual.* From an instrumental perspective, these techniques can improve an individual's efficiency. For example,

---

campaign focused on recycling practices on the UCLA campus. Project participants used hone cameras to document and tag incidences of recyclable materials thrown into garbage cans on campus. The incidents were tallied, mapped, and reported to campus facilities to suggest the most urgent places to add new recycling bins. The goal here is to make clear that CENS PS projects can easily extend beyond self-surveillance.

[13] We want to highlight that self-surveillance data are rarely used only in isolation. More complete self-knowledge often turns on some form of analysis that examines trends and makes comparisons with reference populations, or "peers."

by collecting data about how we spend our time, we can use this scarce resource more productively. In addition to decreasing waste, we can also become more instrumentally efficacious in reaching our goals. For example, if the goal is to watch what we eat, a systematic record of our eating behavior as compared to casual memory can provide a more accurate caloric and nutritional breakdown of the food we consume. As another example, if we are concerned about the carbon footprint we impose on the environment, again self-surveillance of our energy consumption can tell us what kind of emissions we should account for.

[19] There may also be less instrumental and more intrinsic values for the individual. For example, we may have deeply inaccurate (and often self-serving portraits) of ourselves. Self-surveillance may demonstrate, for instance, that we navigate far less an ethnically diverse neighborhood than we suppose; that we waste more energy than our hybrid-bumper stickers signal; that we yell at our children embarrassingly often; that we have implicit biases that we explicitly reject.[14] But precise, data-driven self-measurements, alloyed with legible interfaces (with telling visuals) can force us to confront a more accurate self-understanding.

[20] Benjamin Franklin pursued this sort of self-surveillance to inculcate personal virtue, albeit using low-tech tools.[15] For most of his life, Franklin carried with him a little "account book" recording his daily performance on thirteen separate virtues. When Franklin was 80 years-old, Pierre Jean Georges Cabanis saw the book and remarked:

> We have had in our hands this precious little book. One perceives in it a sort of chronological history of Franklin's mind and character. One sees him develop, fortify and mold all the actions which constitute spiritual perfection,

[14] *See, e.g.,* ProjectImplicit.org.

[15] Franklin never completed writing his planned work The Art of Virtue, but he refers to his practice in his Autobiography. *See* Norman S. Fiering, *Benjamin Franklin and the Way to Virtue*, 30 AMER. Q 199, 200 (1978).

> and the art of life and virtue taught in the same manner as that of playing an instrument or manufacturing weapons.[16]

[21] In Franklin's own words:

> "I was surprised to find myself so much fuller of faults than I had imagined; but I had the satisfaction of seeing them diminish." He also wrote: "On the whole, tho' I never arrived at the perfection I had been so ambitious of obtaining, but fell far short of it, yet I was, by the endeavour, a better and happier man than I otherwise should have been had I not attempted."[17]

Few have the self-discipline reflected in Franklin's subjective, manual recording habits. But new technologies such as Participatory Sensing can automate much of the recording process.

[22] *To society*. The data collected from self-surveillance can also benefit society. Again, from an instrumental perspective, it's not only the self-interested individual who seeks to better herself. A well-functioning society seeks similar ends. This is why CENS has encountered immense interest from the fields of public health, epidemiology, urban planning, and resource monitoring. For instance, as a matter of fighting childhood obesity, it may be crucial to get accurate data about physical activity, food consumed, and exposure to "fast food" advertisements and chains. Self-reports based on faulty memory can provide poor quality data that can be supplemented, improved, or replaced by mobile Participatory Sensing data. Such data could produce better diagnoses and more effective interventions. It also gives data collectors a chance to engage with researchers and policy makers around questions of when, how, and why they might share and learn from their personal data collections. And, from a less instrumental perspective, we recognize that collecting data about ourselves and sharing them with our neighborhoods, groups, and communities can promote a deeper collective self-understanding, not only of the present but also its relation to the

---

[16] (as found in id. at 215-16).

[17] http://www.indianamasoniclibrary.com/articles/tifm/v37/BenFranklinsVirtues.html

past.[18] Indeed, data can become a sort of currency with which we can participate in and help construct communities of memory.[19]

[23] Notwithstanding all these substantial individual and collective benefits, an individual may choose not to engage in self-surveillance because of privacy fears. Fears that such telling data might fall into the wrong hands,[20] be used in unsavory ways, or come back to harm the individual can discourage individuals from collecting the data in the first place. We must therefore confront the oxymoronic problem of self-surveillance privacy.[21]

II. Threshold Clarifications

A. Privacy Metrics

[24] Whenever we confront new information technologies and practices, it's easy and commonplace to raise privacy fears with vaguely Orwellian and Luddite overtones. But a systematic analysis requires, first, some attempt at definitions. What is "privacy," and how might we measure it?

[25] *A standard metric: Control.* The privacy literature typically defines information privacy as the degree to which[22] an individual can control the collection, disclosure, and use of personal data. In other words, privacy is a measure of an individual's power over the processing of information about herself. As shorthand,

---

[18] *See* Kang & Cuff, Pervasive Computing (transparency discussion).

[19] *Cf.* intimacy/friendship theories of personal data disclosure (Fried). Or McKemmish, S. (1996). Evidence of Me.... *Archives & Manuscripts*, *24*(1), 28-45; Appadurai, A. (2003). Archive and aspiration. In *Information is Alive* (pp. 14-25). Rotterdam: V2_Publishing/NAI Publishers.

[20] In addition to hackers, one could worry about underpoliced employees with access to servers. See, e.g., Phil Wong, Conversations about Internet #5: Anonymous Facebook Employee, Jan. 11, 2010 (explaining how employees casually viewed private user profiles with a master word that was a variant of "ChuckNorris").

[21] For discussion of why we chose this term, see *supra* note 5.

[22] Sometimes, privacy is phrased not as a measure of capacity but as a "right" to control the flow of personal data.

we call this the *control* conception of privacy: The more control (of personal data), the more privacy.[23]

[26] This control conception can produce peculiar results if and when someone consents to surrender that control. For instance, if Johnny decides voluntarily to strip in front of a webcam with the intention to disclose publicly his personal data in the form of naked images for all to see and share without any restraints, he is arguably basking in full privacy.[24] That is because the control conception focuses on only the *existence* of control over personal data-not *how* one specifically exercises that control at some given moment in time. Accordingly, a person who successfully makes his data secluded, confidential, and unknown has no more privacy (in the sense of control) than another who surrenders that control and gladly makes his data available to all, on Flickr, YouTube, Facebook, and Twitter.

[27] *An alternative metric: Flow.* As an alternative, one could define privacy not in terms of an individual's control over personal data, but in more macro terms that simply describe the *flow* of categories of personal data within the information ecosystem. In other words, for any particular type of information (e.g., public record data, medical data, or e-mail contents), one could ask where, how quickly, and with what bandwidth does such information flow, either through push/broadcast or pull/search pathways? Under such a flow conception, public record data about ourselves, such as whether we voted, flows faster than medical data, which is treated confidentially by law and custom. Put another way, we have less privacy over public record data compared to medical data.[25]

---

[23] This approach is standard in the legal and policy literatures. Westin, Privacy and Freedom 1967. Other sources. See Nissenabaum 70-71. Solove. Rosen.

[24] Various commentators have observed the weirdness of this result. *See, e.g.,* Anita Allen.

[25] The flow conception has affinities with approaches that define privacy in terms of "constraint on access". See, e.g. Ruth Gavison ("privacy is a condition that is measured in terms of the degree of access others have to use through information, attention, and proximity).

[28] This flow conception is less focused on the specific individual (and her particular decisions about personal data) and more on the type of personal data and how it generally tends to move, as gauged in probabilistic and macro terms, within some information environment. It can, for example, come to a sharply different measure of privacy for webcam images. If we as a society become sufficiently exhibitionist such that most of us regularly and voluntarily broadcast naked pictures of ourselves on the Internet, the flow of such information will have increased and conversely privacy (in the flow sense) decreased.[26]

---

[26] Our conception also has connections to Helen Nissenbaum's approach to privacy which insists that information flows respect contextual integrity. *See* Nissenbaum, H. (2009). *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: Stanford Law Books. Nissenbaum considers privacy to be violated when context-relative informational norms are breached without adequate justification. *See id.* at 140. These norms, in turn, can be identified and understood by analyzing various aspects of information flows including their contexts, actors, attributes, and transmission principles. *See id.* at 150. In this nuanced model, an individual's "control" over personal data is not the sole element in deciding whether privacy has been respected, which feature is similar to the flow approach we offer.

There are, however, differences. Nissenbaum's theory of privacy is ambitious in scope. In particular, she seeks to provide both a descriptive and an augmented normative account of privacy. By contrast, we mean intentionally to be more modest and offer no normative account. Moreover, our use of "flow" is meant to be a simpler metric, operationalized closer to the ground, more amenable to mechanical forms of measurement than violation of a "context-relative informational norm."

To see how our metric differs with Nissenbaum's, suppose that culture changes slowly and incrementally such that most people regularly publicly disclose their GPS trails. In other words, it becomes no big cultural deal. Then, by definition, information would not be flowing beyond expected social contexts. Accordingly, Nissenbaum's contextual integrity might well be preserved and privacy wouldn't be violated, undermined, or decreased. By contrast, according to our flow conception, it doesn't matter that an individual consents or that a culture finds some personal data practice banal: The GPS info is moving more quickly throughout the information ecosystem, which means privacy over GPS data has indeed decreased.

[29] One reason we might care about such privacy loss is if it has predictable consequences, such as gender-differentiated harms in the employment context[27] when past behaviors that seem now to reflect poor judgment are made available for the world to search for and peruse. That would be the case regardless of whether individuals chose, in the wisdom of their youth, to go wild willingly.

[30] To take another example, consider the increasing use of full body scanners at airports known as "backscatter."[28] One could argue back and forth to what extent an individual passenger "chose" to take the scan (in exchange for air travel or avoiding a body cavity search or just making a plane after arriving late to the airport). But from a flow perspective, the "choice" to surrender control over a bodily image is largely irrelevant. The flow of these images will increase with the adoption of these technologies. They are likely to leak out for public consumption, especially when a little photoshopping will place flesh and facial features on top of the scan.

[31] Interesting questions arise from comparing the standard *control* versus the alternative *flow* metric of privacy, and we mean to plant a scholarly flag to mark further inquiry. But those questions are mostly beyond the scope of this Article. Instead, we offer both conceptions as plausible metrics by which we can understand and measure privacy. More important, our case for the Personal Data Guardian (PDG), which we detail below, does not strictly depend on which metric one prefers.

B. Privacy Displacement

[32] Having settled on plausible measures of privacy, we turn to the next threshold question of why we should care about privacy

---

[27] Naked pictures of Jenny, a female associate at a law firm, may impact her career differently than naked pictures of Johnny, a male associate, again regardless of the fact that twenty years earlier both individuals happily exercised their "control" to be filmed naked. See also Anupam Chander's essay (making similar sexism point).

[28] Useful background information can be found at EPIC.org. *See* ‹ http://epic.org/privacy/airtravel/backscatter/› (last visited December 21, 2010).

in the first place. After all, if there's no good normative reason, then any claim that self-surveillance undermines privacy should prompt a collective yawn. To answer this question thoroughly, we need a comprehensive parsing of the values and counter-values served by increasing or decreasing privacy. For example, if we adopt the control conception of privacy, the question for consequentialists would be whether the benefits of increasing individual control over personal data (e.g., encouraging personal experimentation) outweigh the costs of doing the same (e.g., increasing deviant behavior).

[33] As important as such philosophical analysis is,[29] that is neither our comparative advantage nor mission. We seek to avoid much of the normative conversation-but in a transparent way. Our assumptions are these: the current level of privacy (however measured) is normatively tolerable even if not ideal. However, the advent of self-surveillance will materially decrease the amount of privacy in the future, holding all other variables constant. That negative privacy displacement can and should be countered such that privacy later is more approximately the same as privacy now.[30] Again, we are not making the normative case for preserving the status quo amount of privacy from first principles; we're just pronouncing our Whiggish belief. To sum up, we are claiming descriptively that the rise of self-surveillance technologies and practices will decrease privacy and are assuming normatively that that's a bad thing.

[34] *Descriptive claim*. It should not be controversial to suggest that, as a descriptive matter, the advent of self-surveillance will decrease privacy across all if not most plausible measures. After all, engaging in self-surveillance means that highly granular Personal Data Streams will be collected. That data then will be

[29] The literature is sizeable on such matters.

[30] For some, a more symbolic restatement might help. The following isn't actually math, but we provide it just in case it's helpful for some readers. Let p be a privacy function; $p(t_0)$ = privacy at time zero (i.e. right now); $p(t_1)$ = privacy at some future time, t_1. Our prediction is that $p(t_1) < p(t_0)$ because of self-surveillance, holding all other variables constant. We further assume that $p(t_1)$ is less normatively attractive than $p(t_0)$. The goal then is to adopt whatever strategies that will make $p(t_1) \approx p(t_0)$.

uploaded into the "cloud" since information systems now regularly shunt off data onto remote servers, to achieve robustness and flexibility. Increasingly, people have been sharing that data with others in social media sites, rarely with full comprehension of who can access what. And as such practice becomes more popular, routinized, and expected, both social norms and network effects will materially increase an individual's opportunity cost of maintaining her current level of privacy.[31]

[35] Let's return to the CENS Participatory Sensing (PS) case study. Imagine PS being hosted not by a non-profit university but by a private sector firm. As a for-profit venture, this firm has greater incentives to monetize this data in some way, constrained by existing privacy laws and any public relations blowback if their deeds trigger media attention.[32] Monetization often means parsing that data for behavioral targeting and advertising, in ways that the average user is unaware. And if and when those data are shared with third parties,[33] the individual will have even greater difficulty exercising subsequent control over how those data flow. In the end, the individual's power over her personal data will hardly be plenary; also, the flow of that data will have increased.

[36] Now, one could respond that self-surveillance could not possibly decrease privacy because the Personal Data Streams will be uploaded pursuant to the terms of an individual's *contract* with some service provider. Implicitly invoking the control conception of privacy, one could argue that by clicking "yes" on some clickwrap or consent page, the individual has by definition exercised her power to grant to the private firm the permission to do what it seeks to do with her self-surveillance data. Put

---

[31] A few examples might help explain. For those above the age of 40, having a Facebook account may seem entirely optional. For those who are in their twenties, this is much less so.

[32] See, e.g., facebook snafus--beacon

[33] It's important to recognize that one of those third parties might be the government, in some law enforcement or national security project. The state can purchase personal data in the marketplace, subpoena it through legal process, or lean hard on private actors in gray cases to get data.

another way, privacy will not decrease in the future because the collection and processing of any and all personal data streams will have been implicitly or explicitly consented to.

[37] This formalistic objection fails for various reasons. At the outset, this objection is only plausible under the control conception of privacy. By contrast, under the flow conception, the fact that people "agreed" to let firms process their Personal Data Streams would not be relevant to whether privacy had in fact decreased.[34] But even adopting the control conception of privacy, we know that individuals operate under substantial informational and cognitive limitations.[35] Individuals lack perfect information and suffer from information asymmetry about how their data will be used. Individuals make probability calculation errors, and sometimes underweight harms that are low in salience and diffusely distributed. Individuals suffer from regret, which can be understood as a form of intra-personal collective action problem.[36] At the level of market structure, there may be insufficient competition, bundling of products and services, lock-in and switching costs, etc. all of which contribute to the fact that "control" is exercised only formally.[37]

[38] *Normative assumption*. We also believe that many readers will share our normative assumption that we should counter the negative displacement in privacy caused by self-surveillance. But to repeat, we attempt no moral, philosophical, economic, or policy argument in favor of this normative position.[38] If the reader believes, to the contrary, that there is too much privacy now, then the rise of self-surveillance may cause cheer, not concern.

C. Distinguishing Harder Privacy Problems

[39] Our final threshold clarification is to point out that focusing on the domain of self-surveillance, we carve out an easier privacy

---

[34] We suggest, although do not attempt to prove, that this is one reason why the control definition of privacy is faulty.

[35] See, e.g., Paul Schwartz; Michael Froomkin.

[36] *See* Sunstein.

[37] See, e.g., Julie Cohen.

[38] For some such normative account, see Nissenbaum 162-64 (describing both virtues and limits of conservatism in privacy context).

problem than those raised by other pervasive computing technologies.[39]

1. Not Third-party Surveillance of Us

[40] First, the problem is *self*-surveillance, not *third-party* surveillance of us. In the standard privacy problem, personal data are collected by some counterparty in the course of an individual's interaction with that party, typically in some public or quasi-public sphere.[40] For example, a brick-and-mortar store collects your image on a video camera as you walk through its aisles, or some electronic merchant collects information about your browsing and purchase habits as you shop online.

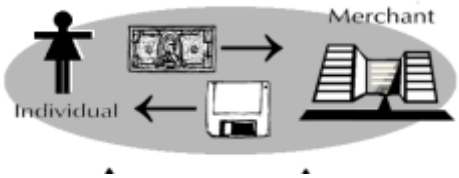

[41] Because the counterparty (e.g., the merchant) collects the personal data in the course of interacting (often executing some transaction) with the individual, that counterparty has some plausible claim to the collected information. For instance, because the personal data were collected through the efforts of the counterparty, it often claims to "own" the data in some way.[41] Given such plausible claims, limiting what the counterparty can do with the personal data once collected raises difficult questions

[39] For a general discussion of pervasive computing in the law reviews, see Kang & Cuff.

[40] For a an early model of cyberspace transactions that focus on the individual, transacting parties, and transaction facilitators, see Jerry Kang, Information Privacy in Cyberspace Transactions, Stan. L. Rev. at 1223.

[41] Implicit Lockean desert theory. Shiffrin. Given the nonrivalrous nature of information, obviously more than one party can "own" various facts, such as the fact that I bought a red scarf on Tuesday for $79. I possess that fact in my short term memory. So does my friend who went shopping with me. So does Victoria's Secret. It is a sort of joint possession.

sounding in terms of liberty ("It's my data since I collected it!"), efficiency ("Better data allow me to serve my customers more effectively!"), and freedom of expression ("The First Amendment allows me to communicate and process this data!").[42] Thus, the typical privacy problem poses a collision between an individual's claim over personal data and the counterparty's.

[42] With self-surveillance, however, the counterparty's interest disappears because the counterparty does not exist. Self-surveillance data are not incidentally created and collected when an individual transacts with some counterparty in the public or quasi-public sphere. Rather, these data are created by purposeful self-initiated surveillance through sensors within the individual's control. Indeed, as a practical matter, these personal data could not be readily collected *but for* the individual's intentionally participating in self-surveillance.[43] Accordingly, no counterparty (e.g. the merchant in our prior examples) has proprietary claim to such data; it didn't collect the data in the first place and often couldn't (under given technological, legal, and financial constraints) even if it sought to.[44]

2. Not Our Surveillance of Third-parties

[43] At the same time, self-surveillance is not our surveilling third-parties. To clarify this point, it is useful to distinguish self-surveillance from other pervasive computing technologies such as

---

[42] Volokh's 1st Amdt. privacy piece.

[43] This is not always the case. See infra Part III.C.2.

[44] There could be gray areas. For example, what if a third party provides an "app" to an individual to engage in self-surveillance. But that "app" has terms-of-service that give the third party some proprietary claim to the self-surveillance data. In this context, via a clickwrap contract, the individual has arguably given some proprietary claim to a third party in exchange for self-surveillance assistance. This muddies the sharper distinctions we drew above, which presumed that no such assistance was needed. We concede that contracting away rights to data can always complicate the tidy picture. In some sense, the infrastructure we recommend below, in the form of Personal Data Guardians and Personal Data Vaults, is designed to obviate such contracts, such that persons can engage in self-surveillance without significant privacy loss.

Lifelogs.[45] A Lifelog is an attempt to produce a complete multimedia record of one's entire sensory experience for permanent personal archive.[46] Imagine having a video camera on your forehead recording everything you see and hear every second of your waking life.[47] Although it has been characterized as a form of "sousveillance" (in contrast to "surveillance"),[48] a Lifelog

[45] For examples of such ventures, see Microsoft's MyLifeBits and USC's Total Recall.

[46] *See* Martin Dodge & Rob Kitchin, *Outlines of a World Coming into Existence: Pervasive Computing and the Ethics of Forgetting*, 34 ENVIR. & PLANNING B: PLANNING & DESIGN 431, 431 (2007) (defining Lifelog as ""unified digital record of the totality of an individual's experiences, captured multimodally through digital sensors and stored permanently as a personal multimedia archive.").

[47] For science fiction iterations, see The Final Cut (Robin Williams).

[48] *See generally* Steve Mann, Jason Nolan, and Barry Wellman, Sousveillance: Inventing and Using Wearable Computing Devices for Data Collection in Surveillance Environments, 1 Surveillance and Socy 331,332 (2003). See Steve Mann, Equiveillance: The Equilibrium between Sur-veillance and Sousveillance 2 (On the Identity Trail, May 2005), online at http://www.idtrail.org/files/Mann,%20 Equiveillance.pdf (visited Jan 12,2008):

Surveillance is derived from French "sur" (above) and "veiller" (to watch). Typically (though not necessarily) surveillance cameras look down from above, both physically (from high poles) as well as hierarchically (bosses watching employees, citizens watching police, cab drivers photographing passengers, and shopkeepers videotaping shoppers). Likewise Sousveillance, derived from French "sous" (below) and "veiller" (to watch), is the art, science, and technologies of "People Looking at". Sousveillance does not immediately concern itself with what the people are looking at, any more than surveillance concerns itself with who or what is doing the looking, Instead, sousveillance typically involves small personcentric imaging technologies, whereas surveillance tends to be architecture or envirocentric (cameras in or on the architecture or environment around us). Sousveillance does not necessarily limit itself to citizens photographing police, shoppers photographing shopkeepers, etc., any more than surveillance limits itself along similar lines. For example, one surveillance camera may be pointed at another, just as one person may sousveill another. Sousveillance therefore expands the range of possibilities, without limitation to the possibility of going both ways in an up-down hierarchy.

for anyone besides a hermit will collect the sight and sounds of other identifiable persons. This is not some "bug" of sousveillance version 1.0; it is instead its central "feature." In this crucial sense, self-surveillance differs from sousveillance. The whole point of self-surveillance is to monitor only the self. By contrast, a Lifelog attempts to record everything that our senses perceive in rich multimedia.

[44] To be sure, incidental capture of data about others will take place even within self-surveillance. And information about others could be inferred from another person's self-surveillance data.[49] That said, a qualitative difference remains: to use a modern day example, it's one thing to take pictures of everything and everyone you see every five seconds (and upload them for the world to share) versus recording quantitative notes about how many times you went to the bathroom in a given day. Whereas Lifelogging has substantial (negative) externalities, self-surveillance threatens others' privacy less. Simply put, capturing data *about oneself* (inward gaze) differs in emphasis from capturing data about others *from one's perspective* (outward gaze). Clearly there are cases where data about oneself does imply information about others (such as with whom you ate dinner or engaged in an activity), but the emphasis and resulting data volume differs in the two cases.

III. Personal Data Guardian

[45] So far, we have cleared brush. First, we identified a nascent socio-technological practice of self-surveillance. Second, although these practices will generate great insight and social benefits, they will also decrease the net amount of privacy, conceived of and measured in various plausible ways. Third, we normatively

---

With the miniaturization of cameras into portable electronic devices, such as camera phones, there has been an increased awareness of sousveillance (more than 30,000 articles, references, and citations on the word "sousveillance" alone), and we are ready to see a new industry grow around devices that implement sousveillance, together with a new sousveillance services industry.

[49] For example, if one has a young child, it will be easy enough to infer her location in a morning commute to school from the parent's location.

assume that this net loss is unattractive. If policymakers agree, how might they counter the displacement?

[46] The natural reflex is to suggest new laws targeting self-surveillance and the service providers that enable the practice. But any such direct regulation seems exceedingly unlikely, not to mention hard to target narrowly. Another predictable response is to suggest some technological fix, which typically touts encryption and efficient individual preference-expression. But so-called Privacy Enhancing Technologies by themselves --without supporting structures -- have historically failed. We take a novel structural strategy: We call forth the Personal Data Guardians.

A. Personal Data Guardian

[47] Our strategy is to introduce into the information ecosystem a new species, which functions as a professional intermediary between her individual client and those who would process the client's self-surveillance data. Specifically, we seek to jumpstart the creation of the profession of Personal Data Guardian (PDG), whose principal mission is to maintain a digital storage locker called a Personal Data Vault (PDV). An individual client would upload her Personal Data Stream into that Vault maintained by her Guardian, instead of into some amorphous cloud owned and operated by some faceless third party.

[48] *Role ideology*. The Guardian would embrace a professional identity of expertise and service, as has been done by other professionals such as lawyers, accountants, financial planners, and librarians. Their role ideology would include the core idea of acting as trustworthy confidantes on behalf of their clients (vis-à-vis third party snoops, subpoenas, and government surveillance), zealous advocates who negotiate for best informational terms vis-à-vis third party application service providers (3P-ASPs), and wise counselors to their individual clients about their decisions regarding self-surveillance data.

[49] *Professional self-regulation*. The Guardian would be an individual human being, licensed as a professional by a state self-regulatory body, which would be most easily created by state

statute.[50] This professional Association would adopt minimum standards to enter into the profession, which standards could include infrastructural capacity, as well as technological, legal, and business competence. The Association would also adopt internal model rules of ethical and professional behavior, whose violation could lead to enforcement actions by the disciplinary arm of the Association as well as malpractice suits by clients. Following the analogy with lawyers, Guardians could partner with other Guardians to create a Firm-in a general partnership or in a Limited Liability Partnership corporation.

B. Personal Data Vault

[50] The Guardian would maintain the Personal Data Vault (PDV), a sort of digital safe deposit box for self-surveillance data.[51] It should provide three basic functions: secure storage, user legibility, and selective third-party access.

[51] *Secure storage*. The Personal Data Stream collected through self-surveillance would be securely uploaded for storage in the PDV. PDVs can be large and physically distributed, hosted across multiple servers.[52] Hosted PDVs can provide a level of

[50] The licensing system could happen at the federal level via congressional statute and supervision by some federal agency such as the Federal Trade Commission. That said, professionals are more typically regulated on a state-by-state level. We also think that as a matter of politics it is more likely that a state legislature than Congress could be persuaded to experiment with a Personal Data Guardian model.

[51] For technical details and descriptions, see Mun, M., Hao, S., Mishra, N., Shilton, K., Burke, J. A., Estrin, D., Hansen, M., et al. (2010). Personal Data Vaults: a locus of control for personal data streams. In *ACM CoNext 2010*. Philadelphia, PA: ACM.

[52] Though they need not be *only* large or distributed. By defining a standard set of protocols, individuals might host their own personal data vault, waiving the benefits of a professional data archivist the same way one might choose to represent themselves in court or store cash under their mattress.

secure storage, robustness, and ease of backup that storing data locally could not provide.[53]

[52] *User legibility*. The personal data stored in the PDV belongs to the individual. But what might it mean for the average individual to access her own data, when digital strings of ones and zeros mean nothing to the typical human being? In some sense, that data has to be made legible to the individual, which means that it must be visualized. We believe that legibility should include visualization of the data in the form of basic descriptive and correlational statistics, as well as visualization across the dimensions of space and time. In other words, basic legibility

---

[53] At a minimum, PDVs must include secure storage, methods for managing individual and third party identities, access control, selective sharing, ability to perform some computation within the vault, data management and audibility, data visualization interfaces, and service interfaces to integrate with third parties. The data store should be redundant to prevent data loss, and should track data provenance and log access to the data. It should also track user changes to sharing rules over time. While the PDV can support strong identity that links data to a unique individual (the data owner), it may also support anonymous and pseudonymous sharing by preventing third-parties from cross-referencing multiple streams to determine identity. There are numerous technical issues involved in delinking identity from data. These include authenticating both users and third party applications, separating personally identified information from data streams, and managing when and how user identity is shared with third party applications. PDV designers will need to consider how users are identified to the PDV operator, and how authentication of data captured on a handset is accomplished, so that malicious parties may not send unauthorized data to the PDV. If personally identifiable information (e.g., name, address, etc.) is held by the PDV, it might be kept apart from the data itself to protect against internal meddling by PDV employees or in some cases subpoena. The PDV and compliant third parties can encourage users to participate in services without those services requiring the identity of the user. This might require a service-specific method of utilizing pseudonyms as part of the PDV API.

should allow simple mash-ups with geospatial data through the use of Geographic Information Systems tools. In addition, Guardians should provide some basic facility to represent the data across time, to show time series and trending. We expect this basic legibility standard to evolve over time, as new techniques emerge from the data analysis community, Guardians, or even users themselves.

[53] *Expertise re 3rd party Access*. Individuals, of course, seek more than basic legibility. They desire more refined applications provided by commercial and non-commercial third parties. The Guardian/Vault structure will allow individuals to share their data in flexible ways with these third parties, instead of opting for an all-or-nothing divulgement of all the personal data. Subsets, statistics only (not the raw data), or scrubbed data[54] can all be made available on a case-by-case basis.[55] Owners or guardians could also request that the PDV run computations on their data and forward only the output.[56] This flexible access can be assisted through automation[57] and be audited (as part of good security).[58]

---

[54] Paul Ohm, Probability of Privacy paper.

[55] The PDV defaults to keeping all data private as they arrive. Access control and sharing mechanisms would allow users or PDGs to change these default policies, setting new sharing policies for particular third parties. For users who do not care to set sharing policies on an application-by-application basis, a range of default sharing profiles would ease the logistical burden.

[56] By hosting some computation within the vault and exporting only outputs, users can access detailed and accurate application outputs while protecting detailed personal information. The simplest way to address this challenge is the installation of common computations as built-in libraries to the PDV. Several types of processing are in common use across participatory sensing applications. One example is inferring transportation modes such as walking, running, biking, and driving using accelerometer and GPS data. Another example is the transformation of GPS data to place name, city name, ZIP code, region name, and country name. This approach resolves the issue of running untrusted application code inside the PDV, but the extent of data sharing choices can be limited by built-in libraries.

[57] Selective access could use filters that can share or protect data based upon variables such as time, activity, and nature of third party requests. These could include warning systems. Filters could make it easy for

Moreover, by providing a single, regularized interface for granting access, we avoid multiple learning curves for the user and decrease the chances of inadvertent disclosures or other personal data mistakes.[59] Finally, the Guardian would be expected to exercise her own expertise and make recommendations on behalf of its clients about various personal data sharing strategies,[60] including which third parties to trust.

---

individuals to express data sharing preferences (e.g. share only data collected between 8 and 10 am; share data only when I am driving; share data only with my doctor). Adaptive filters could learn from user data and use anomaly detection to further help users manage the logistical burdens of selective sharing. For instance, an unusual trip to buy a present for a spouse might be flagged by the PDV, prompting the user to deny a third party application access to that single trip.

[58]Users or external auditors should be able to audit an application provider's storage and access practices, and their use of private data to ensure that it abides by published privacy policies. Moreover, applications might provide tools for users to explore their trails: where and when data originated, what processes were performed on that data, and if and when data was shared. Such tools could become complicated because the volume of audit information scales as a user provides more access to their data. For auditing to be effective and to reduce consumer confusion, it may be necessary to have auditing agencies (analogous to credit reporting agencies or rating agencies). Maintaining per-access audits for each user across several PDVs and providing fast analyses of audit trails provide technical research challenges to explore.

[59] The PDV is similar to approaches implemented to support selective sharing of personal health records, such as Microsoft Health Vault and Google Health. However, the PDV collects raw data generated by the user, rather than records generated by doctors or health professionals. Therefore a different set of tools for visualization and interpretation are necessary, and different regulations apply. A PDV could be built into existing approaches such as Microsoft Health Vault and Google Vault, but the current design of these services does not support self-surveillance data.

[60] Choices of how much and what type of data to share will cause complicated tradeoffs regarding the type and accuracy of calculations that can be performed. The PDV's features for helping users to make

[54] Surprisingly, certain 3rd Party Application Service Providers (3P-ASPs) might actually prefer a Guardian/Vault framework to the status quo. Imagine, for example, a 3P-ASP who is a university researcher, who seeks a better understanding of daily commuting practices in the County of Los Angeles, in order to combat air pollution. This 3P-ASP needs to access not only one Personal Data Stream, but hundreds of thousands. But if all these data are locked away in separate Vaults, how can a researcher access them? One possibility is to allow Guardians to answer federated queries.[61] In other words, we can think of individual Vaults connecting together to form a sort of Personal Data Cloud. One could envision the creation of communication protocols and standards that enable Guardians to collaborate in answering aggregate queries made by service providers, such as our hypothetical researcher. The potential value of this Cloud could alter significantly the underlying cost-benefit calculus and make the Guardian/Vault architecture more attractive from various perspectives. In particular, 3P-ASPs might gladly interact solely with Guardians (and not directly with their represented clients) in exchange for the possibility of working easily with a federated cloud of them.

C. Legal Relations

[55] The fundamental relationship between the individual client and the Personal Data Guardian would be that of the common law's principal and agent, which would mean that the Guardian owes fiduciary duties to her client in handling her self-surveillance data. Consistent with this arrangement, three important duties must be respected.

---

these choices will be important to the ability to support application-specific processing.

[61] This is similar an approach suggested by the Common Data Project (http://www.commondataproject.org/), a nonprofit developing a cloud service which would allow third parties to query sensitive personal data without revealing that data.

1. Fiduciary Duties

[56] *Duty of care*. As a faithful agent, the Guardian must demonstrate a minimum competence in terms of safely storing, securing, deleting, analyzing, and presenting (making legible) an individual's personal data.[62] This duty could be enforced through disciplinary action and the standard common law malpractice tort.

[57] *Duty of confidentiality*. Just as a lawyer or accountant may not ordinarily reveal client confidences,[63] the same would be true with the Guardian. This duty of confidentiality could be enforced through disciplinary action, as well as tort[64] or contract actions.

[58] *Duty of loyalty*. As a fiduciary, the Guardian owes a duty of loyalty to the individual client. But conflicts of interest can arise if the Guardian becomes vertically integrated with third-party application service providers (3P-ASPs). In such cases, what is best for the individual client may not be best for the Guardian or the Guardian's Firm, which would profit from the individual's adoption of its own application services. Instead of monitoring for misbehavior, which has historically been difficult in such contexts,[65] an ex ante structural solution would be cleaner. Just as we don't generally allow law firms to provide vertically integrated services and prohibit attorneys from partnering with non-attorneys in multidisciplinary practices,[66] the Guardian would be similarly quarantined from providing application services.[67]

---

[62] See supra discussion re legibility. Similar issues of minimal competence arise in other sectors that require safe keeping of personal data, such as the financial and health care industries.

[63] In the lawyer context, this often includes the very identity of the client.

[64] The tort could be malpractice. In addition, common law courts could recognize a separate cause-of-action for breach of confidentiality.

[65] Cite to AT&T DOJ case of discriminatory interconnection and cross-subsidization.

[66] Model Rule 5.4 prevents nonlawyers from partial ownership of a law firm. It also prevents lawyers from providing multiple services (beyond legal services), such as accounting or health care, from the same office.

[67] Somewhat complicated. Model Rule 5.7 allows ownership through structurally separate arms. (Consider softer option with net

2. Evidentiary Privilege

[59] In addition to the above three fiduciary duties, the self-surveillance data stored in the Vault would be wrapped by an *evidentiary privilege*, similar to the non-commercial trade secret privilege. In other words, none of the data stored in the Vault could be subpoenaed or introduced into any legal proceeding unless the privilege was waived by the individual, or subject to some clearly delimited exception.[68]

[60] Similar to the three duties discussed above, this privilege could be recognized by state judge application and extension of the common law. Some analogies can be found, for instance, in the recognition of a self-evaluation or self-critical analysis privilege in certain states.[69] In the alternative, state legislatures[70] could simply pass a statute creating the privilege, as some have done for medical committee reports.[71]

[61] *The need for the privilege*. This evidentiary privilege provides substantial benefits to individuals engaging in self-surveillance. The rules of discovery in civil litigation allow for tremendous access to self-surveillance data held by third parties. For example, in federal litigation, the Federal Rules of Civil

---

neutrality-like idea) Thus, if this model were to be adopted, current CENS projects such as PEIR would have to be broken apart into two separate entities.

[68] As with most evidentiary privileges, there could be voluntary as well as inadvertent waivers. Also, there could be specifically identified legal exceptions, such as the "crime/fraud" exception to the attorney-client privilege.

[69] Some such privileges have been recognized in the context of medical committee reports, affirmative acion studies, and environmental impact assessments.

[70] We focus on state legislatures because although Federal Rules of Evidence Rule 501 permits the federal courts broad discretion in applying privileges "in the light of reason and experience" within the federal courts, Congress expressly rejected the adoption of any specifically enumerated privileges. (Wright & Graham, §5421). Federal courts generally apply the privilege law of the states in which they sit.

[71] *See, e.g.,* Flanigan 83.

Procedure provide that "(p)arties may obtain discovery regarding any non-privileged matter that is relevant to any party's claim or defense." Furthermore, "(r)elevant information need not be admissible at the trial if the discovery appears reasonably calculated to lead to the discovery of admissible evidence."[72]

[62] Given this broad scope of discovery, parties in divorces, contract disputes, and tort actions can subpoena from third parties self-surveillance data. Consider, for example, the case of *Ledbetter v. Wal-Mart Stores*, in which plaintiffs Heath and Disa Powell sued Wal-Mart for injuries and loss of consortium allegedly suffered due to an electrical accident that occurred while one of the plaintiffs was fixing an electrical system in an Aurora Colorado Wal-Mart.[73] The alleged injuries included "sleep disturbance and anxiety" as well as "fatigue, cognitive inefficiencies and depression" all contributing to claims for direct damages as well as to the claim for loss of consortium.[74]

[63] Wal-Mart issued subpoenas to Myspace, Facebook, and Meetup.com, seeking information and communications stored by these websites that they hoped would refute the plaintiffs' medical diagnoses and cast doubt on the claim of loss of consortium.[75] When third-party websites are served subpoenas, they typically resist--at least mildly. However, all of the "privacy policies" make clear that they will turn over data when lawfully required to do so. Being served a subpoena is part of that lawful process.

---

[72] FRCP 26(b)(1). Although not all states have adopted the Federal Rules, an quick and admittedly incomplete scan of a sampling states with their own rules of civil procedure of did not reveal any substantial variations in the rules of discovery.

[73] *See* First Amended Complaint And Jury Demand, No. 06-cv-01958-WYD-MJW Document 1 Filed January 30, 2007 (D. Colo. 2007) at 16, 60-71.

[74] Ledbetter v. Wal-Mart Stores, Inc. No. 06-cv-01958-WYD-MJW April 21, 2009 WL 1067018 (D.Colo. 2009) at *1

[75] Now, self-authored text and manually uploaded photographs and videos of oneself is in some sense a primitive form of self-surveillance. However, these these social sites could include applications, for example, that include location streams-- which fit squarely into the definition of self-surveillance.

[64] One might believe that specific privacy laws, such as the Stored Communications Act, prevent such disclosure. Indeed, citing this Act, the various websites subpoenaed in the *Ledbetter* case declined Wal-Mart's request for information.[76] However, this just led Wal-Mart to file a motion to compel discovery against the plaintiffs who, according to Wal-Mart, had "possession, custody, or control"[77] over the relevant information because they could grant or deny access to their accounts. Agreeing with this characterization, the court compelled the plaintiffs to grant the social networking websites permission to disclose the requested information to Wal-Mart.

[65] Although currently there are only a few examples of such litigation strategy, it will likely become common practice. And, although there are other protections against overbroad discovery,[78] no one should feel especially safe about self-surveillance data held by third-parties. The results would differ radically if the self-surveillance data were held within a Personal Data Vault, protected by something like a trade-secret privilege.

[66] *The mechanics of the privilege*. To make our analysis concrete, imagine that as part of a Personal Data Guardian initiative, a state legislature creates the following privilege:

> A person has a privilege, which may be claimed by him/her to refuse to disclose and to prevent other persons from disclosing self-surveillance data stored in a Personal

---

[76] Defendant Wal-Mart Stores, Inc.'s Motion To Compel Production Of Content Of Social Networking Sites No. 06-cv-01958-WYD-MJW Document 185 Filed 05/26/2009 (D. Col. 2009) at 2.

[77] FRCP 34(a)(1)(A) states that this requirement applies to "any designated documents or electronically stored information-including writings, drawings, graphs, charts, photographs, sound recording, images, and other data or data compliations-stored in any medium from which information can be obtained either directly or, if necessary, after translation by the responding party into a reasonable usable form".

[78] Some are built into the federal Rules of Civil Procedure. In other contexts, sector specific privacy laws might provide some obstacles although they may similarly be vulnerable to the motion to compel technique described above.

> Data Vault by a licensed Personal Data Guardian, so long as the allowance of the privilege will not tend to conceal fraud, enable criminal activity or otherwise work injustice. When disclosure is directed, the judge shall take such protective measure as the interests of the holder of the privilege and of the parties and the furtherance of justice may require.[79]

[67] Evidentiary privileges are categorized as either topical or communicative. In other words, they protect either information on a certain subject matter (e.g. trial preparation materials)[80] or confidential communications between two people (e.g., attorney-client). The self-surveillance privilege is designed to be the former, not the latter type since we are interested in protecting the underlying observations collected through self-surveillance and not just the confidential communication[81] between, say, the client and the Personal Data Guardian.

[68] Topical privileges, including this one, might seem overbroad because they are not constrained only to the confidential communications between two select parties. But notice that this privilege is sharply demarcated in two ways. First, as a matter of scope, the privilege only protects "self-surveillance data" that is stored in a "Personal Data Vault" maintained by a licensed "Personal Data Guardian." Second, as a matter of strength, the privilege is qualified. We unpack each of these points.

[69] *Scope*. As a bright line rule, only self-surveillance data that are in the custody of a licensed Personal Data Guardian within a Personal Data Vault enjoys the topical privilege. Accordingly, if

---

[79] (Modeled after Federal Rules of Evidence Rejected Rule 508 - Trade Secrets Privilege)

[80] *See* FRCP 26(b)(3) (also known as "work product").

[81] Unlike the topical privileges that protect facts, a confidential communications privilege applies only to communications. (According to Imwinkelried, "communications" include "expressive statements and acts." A statement or act is "expressive if the speaker or writer subjectively intends the statement to convey meaning to a person such as a hearer or reader. State of mind must also exist at the time of the transfer. A transmission of a preexisting document does not qualify as a communication with respect to privilege." Imwinkelried, 731.)

an individual simply recorded herself and kept the data on her own computer, it would not benefit from the privilege because it is not being held by a Personal Data Guardian.[82] Critics may challenge this sharp limitation of the privilege: after all, if the goal is to protect a sort of information, why should it matter who happens to be holding it? This is a fair point, but we advocate a bright line rule to discourage overbroad assertions of the privilege. When an individual claims the privilege, she may be inclined to do so self-servingly. By interjecting a Guardian as an intermediary, who has professional responsibilities, the privilege is less likely to be abused.

[70] *Qualified not Absolute*. In addition, this topical privilege is not absolute. The proposed statute states explicitly that the privilege may not be deployed to "conceal fraud, enable criminal activity or otherwise work injustice."[83] Further, as characteristic of qualified privileges, every attempt to establish a self-surveillance data privilege would need to pass a case-by-case balancing test at the discretion of the trial judge.[84] The privilege would not give way just because some of the self-surveillance data is "generally relevant" to a party's case or claim.[85] As such, it would effectively stop discovery requests that reflect bad faith, maliciousness, or unnecessary prying.[86] Rather, the judge would override the privilege only if the self-surveillance data are "directly relevant to a material element of the cause of action (or defense) and necessary because the party opposing the claim of

---

[82] We could envision allowing local backup copies of the PDV, for example on separate hard drive, as long as it remains within the networked "custody" of the PDG.

[83] In addition, all privileges can be waived, expressly and inadvertently.

[84] Wright & Graham , at 384. The "judge has discretion to override the privilege claim when the interests it serves to protect are outweighed by some countervailing interest." Wright and Graham, 289). Factors that the trial judge would weigh include: dangers of abuse, good faith, adequacy of protective measures and availability of other means of proof. (Wright and Graham, at 283).

[85] *See* Rutter California Practice Guide: Civil Trials and Evidence, Ch. 8E-A(13)(b)(2).

[86] Wright and Graham, 386. Think prying after divorce, political rivals, insurance or worker's compensation claims.

privilege would be unfairly disadvantaged in proving its case absent access to the" self-surveillance data.[87] And when they do so, judges would take care to use techniques, such as *in camara* review and protective orders to limit public disclosure of the evidence.[88]

* * *

[71] We have provided only a cursory sketch, but the goal of this Article is to suggest the basic innovation, not to provide implementation specifications. Even with this preliminary understanding, we can see how introducing the Guardian/Vault will allow self-surveillance to take place with its attendant benefits while decreasing privacy losses, however measured. For example, under the control conception of privacy, having an expert and loyal agent surely increases an individual's actual (as opposed to purely formal) control over personal data. Consider by analogy a similar relationship in the context of medicine. Having an expert and loyal doctor surely increases our control (actual autonomy) over our own bodies. We reach a similar conclusion with the flow metric of privacy.[89] By role ideology, a Guardian is

---

[87] *See* Rutter California Practice Guide: Civil Trials and Evidence, Ch. 8E-A(13)(b)(2).

[88] Rutter California Practice Guide: Civil Trials and Evidence, Ch. 8E-A(13)(b)(2). Mueller and Kirkpatrick explain that "(a)mong the more common protective measures are orders that the information be disclosed under seal, and that it not be filed in court unless necessary in connection with discovery or substantive motions . . . that (the evidence be disclosed) only to the attorney for the party seeking discovery . . . (and c)ourts may conduct *in camera* inspection." Mueller and Kirkpatrick §5.49.

[89] At various moments, we've raised the domain of medical data. One might naturally wonder whether certain medical privacy laws might apply to the Personal Data Vault. In particular, the Health Insurance Portability and Accountability Act ("HIPAA") Privacy Rule, 45 C.F.R. pts. 160, 164 (2008), "protects the privacy of individually identifiable health information." U.S. Department of Health and Human Services, Health Information Privacy, http://www.hhs.gov/ocr/privacy/index.html (last visited July 14, 2009). But it applies only to such information held by three types of entities:

health plans; health care clearinghouses; and certain health care providers. *See* 45 C.F.R. §§ 160.102(a), 160.103 (2008).

Under the HIPAA Privacy Rule a health plan is "an individual or group plan that provides, or pays the cost of, medical care," and a health care provider is "a provider of services (citation omitted), a provider of medical or health services (citation omitted), and any other person or organization who furnishes, bills, or is paid for health care in the normal course of business." 45 C.F.R. § 160.103 (2008). Neither of these definitions applies to the PDV system, which merely stores (and transmits) self-analytic data.

> A storage-only PDV also does not qualify as a "health care clearing house," defined by the HIPAA Privacy Rule as a:
> public or private entity, including a billing service, repricing company, community health management information system or community health information system, and "value-added" networks and switches, that does either of the following functions:
> (1) Processes or facilitates the processing of health information received from another entity in a nonstandard format or containing nonstandard data content into standard data elements or a standard transaction.
> (2) Receives a standard transaction from another entity and processes or facilitates the processing of health information into nonstandard format or nonstandard data content for the receiving entity. Id.

Because the PDV stores but does not process data nor acts as an intermediary between health plans and health care providers, it cannot be classified under the rule as a health care clearinghouse. Were the PDV system to process raw self-analytic data into a standardized format for the benefit of the user, the HIPAA Privacy Rule would still not apply because the PDV system is receiving the data from the individual and the individual determines where this data is transferred.

A system somewhat analogous to the PDV system, and which is also not regulated by HIPAA, is Google Health (http://www.google.com/health), a free service designed by Google to "store and manage all of your health information in one central (online) place." About Google Health, http://www.google.com/intl/en-US/health/about/index.html (last visited July 14, 2009). The PDV and Google Health systems are similar to the extent that both may store self-analytic health care information and both serve individuals.

invested in slowing down - not speeding up - the flow of personal data. Professional ideology as well as fiduciary law require her to put her clients' interests above that of 3rd parties. Moreover, the Guardian is structurally conflicted out of adjacent vertical markets, which decreases the chance that financial self-interest will warp recommendations.

IV. OBJECTIONS

A. Implausible?

[72] One might believe that the Guardian and Vault proposal is purely academic since no conceivable business case exists for them, and thus, it will never materialize in the marketplace. Put another way, individual consumers will not be willing to purchase their services at a price that would make it worthwhile for the Guardians to enter the profession. Instead, individuals will interact directly with third parties (3P-ASPs), as they do now.

[73] Any prediction in law reviews about whether an entire line of business is economically viable will be speculative. That said, some rough comparisons can provide useful information. For instance, consider what various software and storage services cost in 2010. In terms of secure storage, mozy.com offers unlimited personal backup storage for approximately $50 per year. In terms of privacy-promoting services, web anonymizer proxies such as Anonymizer.com charge $70 per year.[90] Identity theft

---

(However, unlike the PDV system, Google Health's specific purpose is to store all health care information, and it does not store non-health related self-analytics. In contrast, the PDV system is designed to store all self-analytics and exclude non-self-analytic health care information.) On its site, Google Health states: "Google Health is not regulated by (HIPAA) . . . because Google does not store data on behalf of health care providers. Instead, our primary relationship is with the user." Google Health and HIPAA, http://www.google.com/intl/en-US/health/about/privacy.html (last visited July 14, 2009).

[90] discuss cash gift card industry? Some services are free. Onion routers / Tor.

protection services, such as Lifelock, (claim to) guard against identity theft and assist clients who are victimized at $120 per year. Wells Fargo offers a "vsafe" account, advertised as "your personal online safe,"[91] that allows storage of a 1GB of data for $60 per year. It seems plausible, then, that a Guardian could offer basic Vault services to individual clients at approximately $100 per year, which we believe would be inexpensive enough for many individuals to sign up.

[74] What's the value-added for that price? Already, we throw data up to the "cloud." Our genetic information might sit with some genome sequencing company, our time data here, our GPS data there-with nothing but generic privacy statements on web pages and clickwrap licenses. Think how much more comfortable many of us would be if all such data were as safe as if they had been communicated to a personal lawyer in the context of seeking legal advice, and thus protected by something as robust as the attorney-client privilege. Of course, no technological or legal safeguard is foolproof. Even the attorney-client privilege has numerous exceptions, and malpractice actions against lawyers who breach a duty of confidentiality are thoroughly burdensome. But our analysis should always ask: "Compared to what?" Having some protection is better than none.[92]

---

[91] *See* ‹https://www.wellsfargo.com/wfonline/wellsfargovsafe/index› ("The new Wells Fargo vSafe service offers secure online storage for you to safeguard, organize, and access electronic copies of important documents-from birth certificates and immunization records to wills and treasured photos").

[92] A skeptic might say that if it's an Internet data service that doesn't provide immediate gratification (e.g., music, games, pornography), customers won't pay for it, and instead insist on free services financed by advertisements. Obviously, we are not sanguine about the idea of a PDG delivering ads to her clients. One solution might be to sell clients a physical object, such as a hard drive, on the assumption that consumers are more willing to pay for such items. One could imagine Personal Data Guardians selling hard drives that offer local encrypted backup of their Personal Data Vaults. As self-surveillance data are streamed to the PDG, they could be sent back down to a specifically authenticated drive in a reverse-cloud backup. The cost to the PDG of the drive might be $100. But PDGs could sell them to their clients as part of their

[75] In addition to protecting clients from harm, Guardians could teach and advise. After all, the point of self-surveillance is increased self-knowledge. This requires some education, exposure to statistical concepts, and understanding of inferences. Just as the best financial planners help their clients understand concepts such as portfolio diversification, the time value of money, tax deductions (not to be confused with credits), and compound interest, Guardians might do the same for their clients. By this we don't mean personally customized one-to-one tutorials, which probably would be too expensive. Instead, we mean something like the financial literacy training provided by Motley Fool through its website,[93] or the educational materials on the non-profit Privacy Rights Clearinghouse and the Electronic Privacy Information Center.[94]

[76] In this admittedly speculative analysis, we should remember two other moving parts-social norms and the law. First, one could imagine a world where it would seem uncouth, unsafe, and downright shady for a third-party to ask directly an individual for her self-surveillance data. It could be akin to an Internet merchant insisting on your social security number to make a minor purchase. An individual might think to herself: "Why would they do that when a perfectly functioning data vault system exists? What are they trying to do?" And if a corps of Guardians does come into existence, fully embrace their role, and evangelize accordingly, then social norms could emerge strongly against directly depositing self-surveillance data with less trustworthy third parties.

[77] The second moving part is the law. In its bluntest implementation, use of Personal Data Guardians could be mandated in certain circumstances. By way of analogy, in various states, one cannot consummate a real estate transaction without the participation of either real estate agents or lawyers.

---

service for $200, thus producing the $100 mark up necessary to provide their services. This is sheer marketing speculation. We thank Jeff Jonas for conversations about this idea.

[93] See ‹www.fool.com›. The site's trademarked motto is "The Motley Fool: To Educate, Amuse & Enrich".

[94] http://www.privacyrights.org/

In other words, by force of law, an intermediary is injected into a market transaction that makes it impossible or difficult for two parties to interact otherwise. For the most sensitive self-surveillance data (e.g., genetic or medical), this intermediation could be made an immutable legal requirement.

[78] The law could also exercise influence more indirectly, simply by increasing the value proposition of Guardians and their Vaults. For instance, by recognizing an agency relationship that does not exist with typical 3rd parties, the law raises the substantive value of the Guardian-client relationship. In other words, if one wants enforceable duties of care, confidentiality, and loyalty, the best option may be the Personal Data Guardian.

[79] True, it's possible that third parties will offer to do the same through contract. But one would, for example, almost need legal training to distinguish carefully between advertising puffery and actual legal relations. An illustrative example comes from Wells Fargo's vsafe product. In its advertising, Wells Fargo promises safety and security. But in its actual terms of service, Wells Fargo states in fine print: "You acknowledge that by storing copies of your electronic records with us, no fiduciary relationship is created between you and us."[95] Furthermore, no amount of private contracting could replicate the evidentiary privilege we have already discussed. Finally, one would have far more avenues of recourse against an incompetent or disloyal Guardian than a third-party. Besides the contract claim, a client would be able to sue in tort as well as initiate some self-regulatory disciplinary action.[96]

[80] Yet another way to promote uptake is through a combination of pulling strings attached to governmental funding. Academic research institutions that receive government funding through grants and contracts form a significant market for personal data.

---

[95] ¶ 1. Further, there is no evidentiary privilege. The Agreement states "You may understand that we may provide copies of electronic records in your Wells Fargo vsafe Account and our audit logs in response to legal process." ¶ 5.

[96] We don't want to be overoptimistic about self-regulation. We recognize that professional societies in practice serve as only mild deterrents to bad behavior. Abel. But mild is better than nothing.

These institutions are governed by strict national guidelines for the protection of research subjects.[97] Researchers concerned about mandates of respect, beneficence, and justice for research subjects might use Guardians to promote meaningful consent and minimal harm, two tenets of research ethics. Data vaults would allow research subjects to collect study data and then submit that data to participating researchers trusted by vault Guardians. Researchers working on particularly sensitive issues might run federated queries with the vaults, thereby gaining access to aggregate statistics without accessing the raw data themselves. Vaults could help researchers gain approval of their Institutional Review Boards (IRBs) and comply with national guidelines for the protection of research subjects. Incentives for research participants might grow to include funding for PDV subscriptions, much as sensing research incentives currently include access to mobile phones and data plans.

B. Useless?

1. Third-party Surveillance

[81] A second objection is that a Personal Data Guardian is useless because self-surveillance is not the real problem. Instead, the real threat is third-party surveillance of us. Above, we noted that self-surveillance should be distinguished from the harder problem of 3rd party surveillance of us.[98] Although that distinction is crucial, one could say that 3rd party surveillance of us is now so pervasive and detailed that the contents of a Personal Data Vault would not be unique.[99]

---

[97] 45 CFR 46.

[98] *See supra* Part II.C.1.

[99] Here are some other examples. Suppose that instead of getting an individual to spit carefully 5 ml into a sterile test tube, one could get her DNA simply by shaking her hand or collecting the wine glass she's drunk out of. Suppose that instead of placing a software bug that records how we work on our computers, in the near future, everything-browsing, email, calendaring, games-is done through Microsoft or the wireless broadband service provider, who then collects all the information directly.

[82] Take, for example, the CENS Participatory Sensing case study. The core of the Personal Data Stream is location data captured by a GPS sensor voluntarily worn by the individual. But location can be fairly accurately determined through mobile phone triangulation techniques. And soon, even commodity phones will have GPS radios built in. Since location information is then available to the mobile phone provider, such as Verizon, perhaps Verizon's surveillance can produce the same data that self-surveillance would and does produce.

[83] On the one hand, this objection carries much force. If it is true that third parties can collect as much telling data as self-surveillance, then the Personal Data Guardian solution is partial at best. On the other hand, we have good reason to believe that the factual assumption that self-surveillance collects qualitatively more sensitive data than third-party surveillance is now true and will remain so. First, as a technological matter, self-surveillance currently can produce much more telling data than third-party surveillance. Perhaps that gap will narrow as better surveilling technologies go mainstream, but some such gap will likely persist into the foreseeable future.[100] Second, as a political matter, many such technologies at least as deployed by 3rd parties will be constrained since they will be deemed politically and socially unacceptable. Thus, even if the technology could eliminate that gap, laws and norms will likely keep that from happening. And in the meantime, the problem and possibility of self-surveillance privacy remains to be solved.

2. Genie out of the Bottle

[84] Even if an entire profession of Personal Data Guardians comes online, third-party application service providers (3P-ASPs) will have to gain some access to the self-surveillance data in order to provide useful analysis. After all, the Guardians themselves are quarantined out of such services, to avoid conflicts of interest. But this raises the perennial privacy question of what

[100] Consider second and third generation iterations: What if the vault encourages third party surveillers to deposit their data on you into your PDV. So your bank, Verizon, your cable company, your local government, your health record. All your data on you deposited into your account!

to do with third-party transfers. Once the data leave the Vault, won't it in practice lose all protections? This is the genie out of the bottle problem.

[85] This is a serious and difficult problem, which we did not create and is in no way unique to our solution. In fact, the "genie" problem is much worse when personal data are stored by 3rd parties directly, who feel fewer constraints on secondary transfers. By contrast, our Personal Data Guardian approach makes four improvements.

[86] *Parsimony*. In certain circumstance, the Personal Data Stream need not ever leave the Vault. Instead, the third-party could send a script to the Guardian, to run on the data that remains in the Vault. The results could then be sent to the individual, in some legible format. When the data must be released, the Guardian can adopt a parsimony principle that discloses the least amount of data necessary to execute the requested analysis. The Guardian can identify the minimum data type and sampling rate needed by a third party application, and share only that minimum type and amount of data. For example, the Guardian might release GSM cell-tower triangulations rather than more precise GPS data to third-party applications that don't require fine-grained location information. Or the Guardian could release only the amount of time spent driving if actual position is not necessary.[101] Again, this is where the Guardian's expertise and duties to her client can lead to substantial benefits.

[87] *Self-help lockdown*. A second strategy would be to wrap personal data in Privacy Rights Management (PRM) that increases the likelihood that the personal data will be processed only in authorized ways. Similar to the digital rights management (DRM) deployed by copyright holders, PRM could use audit trails[102] to revoke access to user data in the case of violation. Services such as Ephemerizer[103] and Vanish[104] provide

[101] Acknowledge reidentification threats. Ohm paper.

[102] *See also infra* ¶ (89) (discussing TraceAudits).

[103] Perlman, R. (2005). The Ephemerizer: making data disappear (Technical Report No. TR-2005-140). Boston, MA: Sun Microsystems Labs.

[104] http://vanish.cs.washington.edu

plausible examples. It is also possible to imagine that users could take back data if they changed their minds about the consequences of data sharing ("remote revocation").

[88] *Transitivity*. A final strategy would leverage the law of contract. Imagine the Guardian insisting on a sort of contractual transitivity of obligations that flow with the personal data to 3P-ASPs. In other words, before any Guardian would allow a 3P-ASP to access, possess, or process personal data, they must themselves enter into a contract that includes promises by the 3P-ASP to respect the various obligations (confidentiality, care, etc.) that the Guardian has to the client. Moreover, this contract could explicitly list the client as an intended third party beneficiary,[105] with the right to sue the 3P-ASP for breach of its contract with the Guardian.[106]

---

[105] Restatement Second of Contracts § 302. Even without an express designation, in most jurisdictions, intended beneficiary status might still be found since the circumstances suggest that the party to the contract that extracted the promise (the Guardian) did so for the benefit of the third-party (the client). *See* Comment e to Section §304 of the Restatement of Contracts (suggesting that courts look to whether "recognition of the right will further the legitimate expectations of the promisee, make available a simple and convenient procedure for enforcement, or protect the beneficiary in his reasonable reliance on the promise.").

But some jurisdictions, such as New York, resist looking beyond the four corners of the contract. *See, e.g.,* Debary v. Harrah's Operating Co., Inc., 465 F.Supp.2d 250, 263 (S.D.N.Y. 2006). To satisfy these jurisdictions and to avoid the uncertainty that inherently accompanies a judicial determination of intended beneficiary status, explicit identification of the client as an intended beneficiary in the contract itself makes sense.

[106] Courts routinely recognize the right of an intended third-party beneficiary of a contract to recover damages for breach. *See e.g.,* Santa Clara v. Astra USA, Inc., 540 F.3d 1094 (9th Cir. 2008) (local medical clinics allowed to recover damages from pharmaceutical companies resulting from the breach of pricing agreements with the federal government), Vanerian v. Charles L. Pugh Co., Inc., 761 N.W.2d 108 (Mich. Ct. App. 2008) (homeowner permitted to recover damages from a subcontractor that breached a contract with a general contractor to install a floor in her home) Colavito v. New York Organ Donor

[89] Of course, such a legal strategy runs into at least two problems, one technical, the other legal. The technical problem is detection of when 3P-ASPs break their promises. A "TraceAudit" could help. The TraceAudit is a log included in the Personal Data Vault that is meant to increase the visibility of outside access and use of vault data. The TraceAudit requires that third party applications log all activities performed on or with user data. This log is maintained by the Guardian and can be viewed by her clients who may be curious about how their data have been used. It can also be used to detect suspicious events[107] or alert users to possible violations of data use policy.

[90] The legal problem is that even when a violation is detected, what would be the relief granted when damages are hard to calculate? Contract damages are typically limited to unavoidable, certain, and foreseeable economic losses.[108] Unless and until a robust market for self-surveillance data develops (something we are not necessarily eager to see), the violation of contractually transferred obligations would not create any of the standard economic losses for which courts routinely provide compensation.[109] Instead, damages due to breach of privacy terms are more properly considered emotional or psychic losses, forms of harm that courts generally do not recognize as contractual damages unless "the contract or the breach is of such a kind that serious emotional disturbance was a particularly likely result."[110] For certain types of personal data streams, this standard may well be met, and the very fact that the data were stored with the

---

Network, Inc., 438 F.3d 214 (2nd Cir. 2006) (holding that prospective kidney donee had right to sue donor network and others when a kidney that was donated on the condition that he receive it was implanted in another person). By contrast, incidental beneficiaries have no legal recourse in the event of a breach.

[107] Credit card company fraud algorithms.

[108] Restatement 2d of Contracts § 347 and §§ 351-352.

[109] *See*, In re Jetblue Airways Corp. Privacy Litigation, 379 F.Supp.2d 299, 327 (E.D.N.Y. 2005).

[110] Restatement Second of Contracts § 353.

Guardian could help signal that serious emotional disturbance is likely.[111]

[91] To avoid such complications, the best practice would be for Personal Data Guardians to include reasonable liquidated damage clauses.[112] Courts will only enforce provisions that are "reasonable in the light of the anticipated or actual loss caused by the breach."[113] The more difficult it is to determine actual damages, the more latitude courts will grant to those stipulated by the parties.[114] Because the precise level of these damages is difficult if not impossible to quantify precisely, a conservative stipulation of psychic losses should pass judicial scrutiny. The possible threat

---

[111] Whether violation of contracts guaranteeing privacy allow for damages for emotional distress has seemingly turned on the nature of information that was improperly disclosed. See, e.g., Trikas v. Universal Card Services Corp., 351 F.Supp.2d 37, 46 (E.D.N.Y.2005) (no damages for improper disclosure of credit report) compared to Huskey v. National Broadcasting Co. 632 F.Supp. 1282, 1293 (N.D. Ill. 1986) (prisoner allowed to recover damages for improper broadcast of images of him incarcerated on national television). If we assume that this is ultimately grounded in the principles of foreseeability of particular harm by the contracting parties enshrined in Hadley v. Baxendale, 156 Eng.Rep. 145 (Exch. Div. 1854) then the existence of a Personal Data Guardian and the act stipulating damages itself would seem to alert both parties to the harm and allow for recovery of damages.

[112] *See e.g.,* E.E.O.C. v. First Citizens Bank of Billings, 758 F.2d 397, 403 (9th Cir. 1985) (Liquidated damages are compensatory, not punitive in nature), In re CP Holdings, Inc., 332 B.R. 380, 389 (W.D. Mo. 2005) (citing Paragon Group, Inc. v. Ampleman, 878 S.W.2d 878, 880 (Mo.App.1994) ('Liquidated damages are a measure of compensation which, at the time of contracting, the parties agree shall represent damages in case of breach.' Contrarily, penalty clauses represent punishment for breach), U.S. v. American Motorists Ins. Co., 689 F. Supp. 1569, 1572 (Ct. Int'l Trade 1987) ("True liquidated damages are not penalties. They are compensatory in nature, providing a measure of recovery when it appears at the time a contract is made that damages cause by breach will be difficult or impossible to estimate").

[113] Restatement Second of Contracts §356.

[114] Id, comment b.

of class action aggregation of small claims would increase the stakes.

[92] In the end, nothing prevents a 3rd party from adopting these four strategies right now-except that it is not obviously in its financial self-interest to do so. By contrast, Personal Data Guardians would have greater motivations to implement these strategies, as a part of their professional self-conception as well as competitive advantage.

CONCLUSION

[93] In privacy debates, any new problem is often met by calls for direct regulation or laissez-faire trust of the market. Our approach lies very much in between. Instead of direct behavioral regulation or blind faith in the market, our strategy is to modify indirectly the information ecosystem by introducing a new species, the Personal Data Guardian (PDG). This new creature would be a faithful agent to its client and would store self-surveillance data in its Personal Data Vault. The PDG would also act as a professional intermediary with third-parties who seek access to such data.

[94] Although we have painted with broad strokes, we believe that the PDG framework is a viable, concrete solution to the problem of self-surveillance. What's more, if the Guardians come to be, they will themselves become invested stakeholders, who can shape and alter future privacy policies in this and other domains. Indeed, if the framework functions well in this context, it could be expanded organically to help solve adjacent or related privacy problems.[115]

[95] Novel solutions to privacy problems have become scarce. Simple inspection of the privacy landscape demonstrates that industry self-regulation, self-help encryption, and formalistic notice-and-consent clickwraps are not up to the task. The time for the Personal Data Guardians is at hand.

[115] Medical records.