Debate | Whither Privacy?

Paul Rosenzweig

Red Branch Consulting, PLLC and George Washington University, USA. Paul.rosenzweig@redbranchconsulting.com

Is there a future for online privacy?

Or, more pointedly, given that the vast capabilities of data analytics have come to be a lynchpin (if not *the* lynchpin) for much of government counter-terrorism activity, is there any way to square the need for counter-terrorism data surveillance with a sense of data privacy? In other words, can we have our cake and eat it too? Can we gain the national security benefits of dataveillance (to use a phrase that seems to describe the activity) while also protecting civil liberties and privacy?

It seems that the two will be increasingly in tension if we fail to adjust our conceptions of privacy to account for technological change. Just the other day, the European Network and Information Security Agency concluded that even if we wanted to create greater online privacy (through, in that instance, a right to be forgotten) the trend of technology is making that increasingly impossible to achieve—at least in the classic sense that privacy advocates mean, where a user is in complete control over data about himself.

For a number of years I have been writing critically about our current approach to this conundrum. I have none too charitably called the current set of privacy rules that form the bedrock of much of our thinking “antique privacy.” The more I see of the world around us, the more convinced I am that the description is apt.

Indeed, just after I sat down to begin writing this brief article, the Petraeus affair exploded into the American media (perhaps it made less of a splash elsewhere). While, obviously, the titillating aspects of the story drew the most attention, buried within the story was the, to me self-evident, realization that even America’s senior spy chief has little, if any, remaining digital online privacy.

What, then, are we to do about this? There is, I think, a better way. Let me outline it in this brief note.

First, if we are to adequately protect privacy we have to begin to recognize that privacy is an instrumental value, not, in my judgment, an absolute. By this, I mean that we need to look at privacy as a construct we use to protect other important values—things like autonomy, self-determination, democracy, and liberty of conscience.

In the context of national security dataveillance, the underlying value to protect is, I would posit, the value of anonymity—in effect, the ability to walk through the world unexamined. Or, perhaps more realistically, what we mean by anonymity is that even though one’s conduct is examined, routinely and regularly, both

with and without one's knowledge, nothing adverse should happen to you without good cause. In other words, the veil of anonymity that is now so readily pierced by technology must be protected by rules that limit when the piercing may happen as a means of protecting privacy and preventing governmental abuse. To put it more precisely, to my mind the key to reconceiving privacy in a way that effectively responds to the new technical reality lies with seeing privacy's principal virtue as a limitation on consequence. If there are no *unjustified* consequences—that is, consequences that are the product of abuse, or error, or the application of an unwise policy—then, under this vision, there is no adverse effect on a cognizable liberty/privacy interest. In other words, if nobody is there to hear the tree, or identify the actor, it really does not make a sound.

All that is fine in theory, but what does it mean in practice? It seems to me that it means we must now focus a great deal more closely on the means/ends/consequence fit of any proposed dataveillance program. This idea of considering privacy in a retail way is, I think, fundamentally different from the wholesale, *ex ante*, way we think about privacy now. Under this paradigm, the questions to be asked of any dataveillance program begin with questions of its purpose, efficacy and scope: to what end is the dataveillance directed? How close is the fit between the data analytic questions asked and the end in question? Is it voluntary or mandatory?

Once we understand the means contemplated and the intended scope, we can ask a set of questions about the program's effects on data subjects: what is the consequence of identification? What is the trigger for that consequence? Who decides when the trigger is met? etc. These questions are the ones that really matter, and questions of collection limitation or purpose limitation, for example, are, in my judgment, distractions from the main point. I don't say this with joy, but rather in recognition of technological reality. Practitioners who maintain fealty to the existing privacy structures will, inevitably, find themselves in the backwater of policy making.

The right answers to these more particularized, retail questions of consequence will, of course, vary depending on the context of the inquiry. We will have a different evaluation of, say, a broad based airport screening program that uses link analysis to identify aviation security risk than we will of a more narrowly focused inquiry that, say, assists in the rendering of employment decisions in a critical industrial sector.

Perhaps you think this outline of retail consequential privacy analysis is still too theoretical. So to close out this short essay let me work through a few examples as a way of understanding some possible rule-sets for this type of analysis (all with the caution that my thoughts here are relatively tentative and subject to revision). Let's consider four distinct situations where dataveillance might be applied:

- 1) Comprehensive data review as a condition of employment in a critical infrastructure (such as, say, a nuclear facility);
- 2) A similar review, but as a condition of employment in a non-critical sector (such as, say, education);
- 3) Voluntary data review as part of a frequent traveler program, like the Department of Homeland Security's Global Entry/TSA Pre-Check program; and
- 4) Involuntary dataveillance of limited data set (name, DOB and gender) as a condition of airplane boarding, as with TSA's Secure Flight program.

Our first cut at the analysis might look something like the below chart (where I am describing current practice, rather than an idealized vision):

	Nuclear Plant	Teacher	Global Entry	Secure Flight
Voluntary	No—unless you don't want the job	No—unless you don't want the job	Yes	Yes and No. You don't have to fly—but if you choose to do so you must submit to Secure Flight
Goal of Dataveillance	Significant consequences (threats to critical infrastructure) with low probability	High probability but comparatively low consequence event (e.g. teacher molestation of a child)	High consequence (air hijacking or disaster) with low probability	—As adjacent— [Though note that early conceptions of the program included anti-crime goals as well]
Effectiveness/Fit	Tight fit—Narrow criteria and a focused threat. Database sources likely include commercial	Loose fit—Broader criteria and a broader threat Also commercial sources	Medium fit—Broad threat with detailed screening Exclusively government data sources	Loose fit—Broad threat with minimal criteria Also government data sources
Consequence of Identification	Loss of employment	Loss of employment	Secondary screening and/or denial of travel	Secondary screening and/or denial of travel
Trigger	Terrorist record or perhaps any criminal record	Criminal record or, more narrowly, child abuse record	Terrorist connection/watch list or criminal record	Terrorist connection/watch list
Standard of Proof	Negligible—Just a match	Negligible—Just a match	Detailed adjudication	Negligible—Just a match
Decision maker	Employer	Employer	DHS	TSA

So, what can we learn from this exercise? A great deal, I think.

For one thing, it strikes me as anomalous that the decisions that have the greater long-term consequence (e.g. denial of employment) are the ones that have the least significant degree of proof requirement and are adopted through the least formal decision making mechanisms. By contrast, the dataveillance program that is the object of the greatest public scorn (airplane flight screening) is the one where lesser consequences are imposed, fewer data sources consulted and a more formalized process of identification is used. We can speculate that this is because of the universal and involuntary nature of the program but, nonetheless, my preliminary thought (the most that can be offered in this brief note) is that much of our focus and concern for adverse privacy consequences may be misdirected at less significantly intrusive programs.

A second issue that leaps out from this discussion is that the standard of proof and the identity of the decision maker are all, at least for now, tied to non-judicial, administrative proceedings. We might well

consider whether or not there are certain dataveillance consequences that need a more formalized (perhaps *ex parte* judicial) approach.

I do not, in any way, suggest that the questions I have asked are necessarily the correct ones or that they are comprehensive. To the contrary, I suggest them as notional for the purposes of stimulating discussion and analysis. But to my mind they are the right questions to ask about digital privacy today. We would advance online data privacy far more if we reoriented our concerns away from worries about collection limitations and data retention rules toward a more robust understanding of consequential instrumentality.