NSA Management Directive #424: Secrecy and Privacy in the Aftermath of Edward Snowden

George R. Lucas, Jr.

Whatever else one might say concerning the legality, morality, and prudence of his actions, Edward Snowden, the former U.S. National Security Agency (NSA) contractor, is right about the notion of publicity and informed consent, which together constitute the hallmark of democratic public policy. In order to be morally justifiable, any strategy or policy involving the body politic must be one to which it would voluntarily assent when fully informed about it.[1] This, in essence, was Snowden's argument for leaking, in June 2013, the documents that revealed the massive NSA surveillance program:

> So long as there's broad support amongst a people, it can be argued there's a level of legitimacy even to the most invasive and morally wrong program, as it was an informed and willing decision. . . . However, programs that are implemented in secret, out of public oversight, lack that legitimacy, and that's a problem. It also represents a dangerous normalization of "governing in the dark," where decisions with enormous public impact occur without any public input.[2]

What, however, is inherent in being fully informed when it comes to surveillance? Much literature on informed consent dwells on the problematic nature of voluntary consent, given inequalities in power between those asking for that consent and those expected to give it, the uncertain epistemic features of information about the contents of a policy, and the risks associated with pursuing it. Physicians, for example, routinely object that laypersons cannot voluntarily consent to treatment because they are often unable to appreciate the risks and costs (vis-à-vis the benefits) of complicated medical procedures, and are in any case frequently so traumatized by their illness or physical vulnerability as to be willing to consent to almost anything.[3] In other areas of scientific inquiry, as in

Ethics & International Affairs, 28, no. 1 2014, pp. 29–38.

doi:10.1017/S0892679413000488

anthropological field research, informed consent is rejected or resisted on account of a mistaken impression that it involves signing a release form of some kind (what we might term "procedural informed consent"). Even more often, it is resisted over concerns that the subjects of study, once fully informed about it, will view it as an unwarranted intrusion into their cultural privacy and withhold their permission; or, perhaps even worse, that they will cease to "act naturally" and instead start to "perform" for the external observer, thus ruining the authenticity of field data collected.[4]

So, what position should we adopt regarding the prospects of giving collective informed consent to a policy of cyber surveillance whose technological sophistication and scope stagger the imagination, and whose implications neither we nor those implementing the policy can possibly foresee? Not only do "we" (that is, the collective body politic) lack the technological sophistication to fully appreciate all the dimensions involved in the workings of this policy, but, as in the case of anthropological field work, we may need to be deliberately blind to some of the features of our own surveillance, lest we (the protected) and those seeking to harm us (criminals, terrorists, agents of adversarial nations) all begin to alter our respective behaviors in response.

Let me propose an approach embodied in a number of generally accepted and morally justified practices that may offer some insight into our current cyber dilemma. We can draw an analogy to a variety of practices that seem to entail secret and clandestine elements, while at the same time exhibiting respect for individual privacy and confidentiality. Examples include patients being asked to participate in "double-blind" medical experiments, undercover agents working as part of a domestic police effort to apprehend criminals (such as drug dealers) or to prevent the fomenting of criminal conspiracies, tenure deliberations and blind peer review in academic institutions engaged in the pursuit of scientific research, and confidential agreements among heads of state in the pursuit of policies clearly in their respective public's interests. Another, perhaps less unproblematic, example is the practice of espionage agents of rival states to engage in HUMINT (intelligence gathered by interpersonal contact) and SIGINT (intelligence gathered by the interception of "signals") activities, which constitute tolerated practices that fall outside the normal bounds of law and morality, and that are allegedly carried out in the name of state security and defense.

Sadly, even if an appeal to these analogies is successful, it will not undo the grave damage that Snowden's revelations have done to the public's trust in the

workings of the U.S. government—especially when that government has been discovered pursuing secret policies, governed, in turn, by secret laws and courts. Nonetheless, examining these analogies may help in the public discussion that Snowden apparently hoped to encourage concerning (in his own words) policies that might seem "evasive and morally wrong." The common feature of these otherwise disparate analogies is the concept of informed consent. Following these examples, the NSA program of massive surveillance (just like a program of double-blind medical experiments, for instance) could—and, I believe, should—seek to inform, and to obtain voluntary consent from, the subjects of security surveillance, making it clear that there will be certain features of the program *that subjects of a policy designed for their protection will knowingly and voluntarily consent to remain ignorant of* for the sake of the effectiveness of that surveillance. Certain procedural details in all of the purportedly analogous instances remain undisclosed, or "secret," but the general policies themselves (the experiment, undercover police work, big-data surveillance, and so on) are not "secretive." Their *existence* and *operation* are disclosed, and the nature of the risks generally entailed are clear to the affected parties, or to those on whose behalf the policies are being carried out.

What Price Privacy?

If these analogies work, we then need to start over by fully disclosing the generic working features of the program outlined in the heretofore secret memorandum, "NSA Management Directive #424." This document, in sum, contains all that Snowden's stolen and disclosed files have thus far revealed, albeit in bits and pieces. Specifically, this includes PRISM (revealed in June 2013), XKeyscore (revealed in July 2013), and (more recently) the Enterprise Knowledge System and data-chaining program, coupled with the extent and scope of data collection of various types throughout this time. "Mainway," for example, is the code name of the massive data storage unit nearing completion near Bluffdale, Utah, into which the NSA has been collecting or storing some two billion "record events" per day since at least 2010. When fully functional, as Snowden's public disclosure of this memorandum reveals, the capacity of the Utah site will exceed ten times that amount daily (20 billion records). These record events include logs of telephone, Skype, and cell phone calls, tweets, Facebook posts, emails, Internet-site visits, GPS coordinates, and so forth. The collection of all such data is subject

to a variety of fairly complex legal regimes and accountability measures. One might characterize the data collection as the amassing of an enormously large haystack, within which intelligence and law enforcement agencies would, under current U.S. law, be entitled to search for a very few, legally distinct needles.

The Enterprise Knowledge System consists of a variety of relational databases that are legal subsets of this haystack (like PRISM), together with analytical tools and risk-analysis or data-chaining programs (like XKeyscore). Collectively, this system constitutes the means by which the enormous trove of so-called metadata is parsed, subdivided, and finally "mined" or analyzed in accordance with the various legal regimes and presidential directives that govern the permissible use of this information.

"Data-chaining" is the linking up into a variety of patterns of these enormous amounts of otherwise individual and seemingly random, obscure, and even trivial "record events."[5] The result is a kind of topographical mapping of patterns of movement and communication that filters out or excludes the metadata of most of us—unless that data can conceivably be gathered up into a pattern that is interesting . . . or suspicious. The result might be helpfully compared to the kind of "whiteboard" that investigators create to help them solve a crime, connecting the various clues and pieces of data. Except that this whiteboard would encompass a degree of detail and a diversity and quantity of data that would vastly exceed the capacities of human investigators. And, perhaps even more significantly, this data is constantly being assembled prior to, rather than in the aftermath of, the commission of a specific crime. Such activity constitutes the kind of preemptive self-defense that is permitted under domestic law in certain restricted situations to foil criminal conspiracies, but which is outlawed under international law in the case of interstate conflict.

XKeyscore is an analytical program that reviews these connections and assigns a resultant risk-analysis score (vaguely analogous to the quantitative risk analysis offered in a FICO credit score). This score can be calculated in principle for each and every person whose data is incorporated in the database, but more likely and effectively associated with suspicious correlations and connections of data clusters, rather than merely of discrete individuals. That score may ultimately help analysts determine whether there might be probable cause for closer scrutiny of any of the non-random patterns of interpersonal communication discerned.

The very scope and ambition of this effort may elicit immediate and horrifying comparisons with Stasi surveillance programs in East Germany.[6] Certainly, that

comparison was foremost in the public reaction to Snowden's revelations within Germany itself. I do not find these specific comparisons very accurate or convincing, but the advent of cyber surveillance as a form of preemptive national self-defense nonetheless highlights two very critical and as yet poorly understood features of cyber conflict in general that are worthy of attention and concern.

First, in law and ethics we distinguish clearly between domestic law enforcement and its monopoly on the use of force under a rule of law and international armed conflict, during which domestic laws and the rule of law itself are often seriously eroded. A critical feature of the advent of cyber conflict is that it has blurred the distinctions between what were once very different and relatively clearly distinguishable levels of activity and conflict. Serious cybercrime and cyber espionage, however, are increasingly straying into an area in which nations can, with increasing plausibility, declare these to be the equivalent of an armed attack by another state. In particular, such charges can, with increasing plausibility and legal verisimilitude, be lodged against states that willingly harbor or tolerate terrorists or international criminal organizations within their borders.

Second, another, even more important, difference with traditional conflict is that the pursuit of cyber strategy and the employment of cyber weapons and tactics have been largely under the control of intelligence agencies and personnel, whose rules of engagement are vastly different from those of conventional military combatants, or from those of agents of domestic law enforcement. Spies and espionage agents are generally engaged in activities that do not rise to the level of a "threat or use of force" under international law, let alone of armed conflict between states, but many of these activities constitute criminal acts in the domestic jurisdictions within which they take place. Conventional war, by contrast, is understood to occur within zones of combat in which the conventional rule of law has broken down, and to which only the international law of armed conflict applies. This latter legal regime is far more tolerant than domestic law regarding the permission to pursue conflict with deadly force (as Jeff McMahan and David Rodin have rightly pointed out.).[7] The Law of Armed Conflict does, however, encompass certain basic moral principles—noncombatant immunity and proportionality, for example—that do not arise as constraints in the pursuit of espionage. Likewise in domestic law enforcement, unrestricted surveillance, not to mention the use of force, falls under strict legal regimes with accountability, oversight, and at least some degree of transparency, including (most importantly) challenges and adversarial review of the policies and procedures pursued. None of these

international or domestic legal constraints apply in the case of espionage, surveillance, or covert action.

The emergence of relentless and unrestricted cyber conflict over the past decade has seriously eroded all of these firewalls against governmental overreach. Thus far, this fundamental distinction regarding rules of engagement has not been fully acknowledged, let alone well understood. Cyber war is nothing less than unrestricted warfare carried out by spies and espionage agents—that is, by persons who do not think themselves bound by any legal restraints—rather than by conventional combatants trained in the Law of Armed Conflict. Unrestricted warfare is *not* legally permissible or morally justifiable in the conventional case, but it is routine practice among agents of espionage. In cyber conflict, and in planning for it, many of the weapons and tactics are specifically designed to operate against civilians and civilian (noncombatant) targets, a feature that would be illegal, and decidedly immoral, in the conventional case.[8]

Defense of the State and its Citizens vs. Oppression and Political Control

The intent of the state (or, more specifically, its security and intelligence organizations) in conducting such surveillance seems of paramount importance in judging the ethics of surveillance policies. In the present case, we encounter a clear conflict between a hitherto reasonably well-established and accepted norm —that is, privacy—and the alleged benefits that come from violating that norm through the mining of "metadata" by the NSA. Through a peculiarity of the history and evolution of the Internet, much of the world's Internet traffic travels across routes and through switches that lie geographically within the United States. Hence, the United States is in a unique position to monitor and survey not only its own Internet traffic but that of the globe. It is a well-established fact that the denizens of cyberspace, no matter of what nationality, have a strong anarchist and libertarian streak. So the outrage, both domestic and international, over the revelations that the NSA was "monitoring" so much of our electronic communications has been palpable in some quarters. We might take this as strong evidence that the norm of privacy is widely valued from a philosophical perspective as an important right, a kind of exercise of individual autonomy that ought to be strongly protected.[9]

Defenders of the practice of email packet-sniffing and metadata mining more generally, however, reply that there is a second important norm: the security of citizens. That is to say, there may be a legitimate tension or conflict between a longstanding and widely (although *not* universally) accepted norm of privacy—one that functions with special vigor in the cyber domain—with a second, and at least as important norm: namely, that ordinary citizens minding their own business, who have done nothing especially wrong, do not deserve to be unduly subject to grave but avoidable risks of harm. This second norm I will term the "security norm." Security of life, person, and property is one of the most basic and fundamental of human rights; and it is one of the chief functions of just and minimally-rights-respecting governments to ensure that all citizens enjoy it.[10]

The tension between these two norms is hardly new or unfamiliar. It was, after all, one of the U.S. Founding Fathers, Benjamin Franklin, who observed that a people who sacrificed privacy (liberty) for the sake of security would end up with little of either. But Franklin was not referring to security in the deeper sense discussed here. Instead, he was denouncing a cowardly and risk-aversive desire for preserving the existing political status quo at all costs, as opposed to standing up for justice and basic human rights. In our own time, threats to, or violations of, privacy have been strongly resisted not merely as an intrusion upon personal liberty but as an ominous foreboding of increasingly totalitarian or authoritarian control. Hence, the thinking goes, if the United States currently monitors the patterns of email traffic among foreign nationals in an effort to combat terrorism and criminal conspiracies, what is to prevent it from opening and reading those emails, or actually listening to those telephone conversations? And what is to prevent it from moving from voyeuristic intrusion for security purposes to the suppression of all free speech and action for political purposes?

Citizens of a rights-respecting and rule-governed nation might reply that the degree of plausibility and severity of the threat justifies their collective response to it. As with the formation of the Transportation Security Administration in the United States in the aftermath of 9/11, or as in the 1950s civil defense campaigns throughout Europe and North America against the threat of thermonuclear war, citizens would likely be willing to accept considerable inconvenience, and attendant limitations or abrogation of their freedom and privacy, *if* the threat of harm were severe enough—Ben Franklin's warnings notwithstanding. Americans, in particular, might also be well instructed to review the grave threats to liberty and privacy that the U.S. government has already perpetrated in response to

potential or imagined threats (as Americans did with J. Edgar Hoover's FBI wiretapping and McCarthyism).

This is not an impossible puzzle to solve, and responding to a threat perceived to be grave need not degenerate into the kind of fear and hysteria that fostered those abuses of government power in the United States during the 1950s and 1960s. What is required to guard against the hysteria and abuses is, of course, due process and, in particular, adversarial review. Some transparency—that is, admitting that surveillance actions are taking place, justifying them, and clearly stating who is exercising accountability and oversight, if not precisely how it is being exercised—is also desirable to prevent arbitrary abuses of privilege. It is also important to distinguish between types of intentions. For all the denunciation of the NSA's activities by the Chinese, in particular,[11] unlike in China, there is no attempt—or even intent to attempt—on the part of the U.S. government to "lock down" the Internet, control dissent, or otherwise pry into people's personal lives and beliefs in an effort to control their political expression.

It is a serious equivocation to draw a parallel between the past behaviors of corrupt or authoritarian regimes and the future intentions and behaviors of basically democratic and minimally rights-respecting ones. For example, to compare the U.S. actions to those of China regarding data control, as some have done, is completely wrongheaded. The intent of the Chinese government is not to protect individuals from harm, but to control them against their will. The intent of the U.S. government is precisely *not* to control or intrude on individuals' private lives, but to perform its legitimate duty to protect citizens from an unreasonable threat of harm. Given this critical distinction, we might conclude that espionage and surveillance designed to prevent criminal conspiracies or outright warfare, and which aim purely at attribution and denial of anonymity rather than invasions of individual privacy, may be morally permissible if and only if:

- the focus of the surveillance is limited insofar as possible toward legitimate security objectives and military targets, and the sole purpose is to prevent armed attack or its equivalent;
- the harm threatened is genuine and reasonably well-defined, and there exists legally-defined "probable cause" for surveillance;
- there is suitable transparency, oversight, and accountability for the program of surveillance, with full adversarial review of the legal permissions ultimately granted;

- the individual privacy of third parties and bystanders is not invaded, and no harm is done to civilians or their property; and,
- most particularly, the broad scope and reach of surveillance efforts are known to and approved by the public, on behalf of and for the security of whom such surveillance is undertaken (informed consent).

In the case of massive, classified surveillance (quite unlike the case of a relatively limited, double-blind medical experiment), however, this last condition might seem to impose what would prove to be an enormous and infeasible undertaking. No doubt it appeared so to political leaders who gave their consent to the secret policy they pursued. But imagine that the Director of National Intelligence or the Director of the U.S. National Security Agency—in the numerous speeches that both these and other government representatives gave in the years leading up to the Snowden disclosure—had simply divulged the bare outline of the policy. The precise details, including the scope of surveillance and the specific algorithms and programs used in data-mining and detective data-chaining, would not need to have been disclosed. Instead, they could have made a simple, frank statement to American citizens and citizens of allied countries that the intelligence services of these states were cooperating to monitor the patterns of telephone and electronic communication in the cyber domain in an effort to discern and prevent criminal and terrorist conspiracies from coming to fruition. Further, what could have been fully disclosed was a procedure for providing accountability and oversight to the process.

In other respects, the foregoing caveats hardly comprise a complete or sufficient list of guiding precepts, but I do believe they encompass some of the "lessons learned" from the Snowden fiasco. It would be an abuse of the legal permission defined above, for example, for a human "cyber warrior" to listen in to Skype conversations between private adults engaged in no crime, or to read their emails for amusement, let alone to impede their free activity or interfere with their political rights of free expression. That such things were allegedly permitted to occur, and that the general outline of the program itself was not acknowledged by the political establishment, have done serious harm to the public's trust in the security agencies of its government. This can be restored only by adhering in the future to basic principles of sound security as outlined above, including the adherence of administrators of such programs to a basic code of conduct that will serve to guide and limit their actions to justifiable objectives.

NOTES

[1] One interesting and vexing problem is to show how this principle works collectively in a democracy, that is, how it applies to the body politic or general will, as opposed merely to the protections accorded individuals in a liberal state. Is it, for example, a necessary feature of any morally justified rule of law that such a regime requires full transparency and an absence of secrecy, and, in particular, that it mitigates against "clandestine" laws and policies? What about those clandestine activities undertaken on behalf of public security—for example, interstate espionage, but also undercover police work, confidential interstate agreements, and the like? Is "collective voluntary consent" adequately signified, for example, in the substantive content of existing law? Or must there be, in addition, some sort of "public declaration" of how specific legal statutes are generally being interpreted or applied, especially in instances where the resulting policies would without question prove controversial, and the details must, perforce, be concealed?

[2] James Risen, "Snowden Says He Took No Secret Files to Russia," *New York Times*, October 8, 2013.

[3] Most recently, such a controversy has arisen over the informed consent (or lack thereof) in the experimental treatment of prostate cancer during a 1950s study funded by the National Institute of Health at Columbia University using homeless men from the Bowery in New York City.

[4] G. R. Lucas, *Anthropologists in Arms* (Lanham, Md.: AltaMira Press, 2009), e.g., ch. 6.

[5] I choose and flag this designator specifically, since the relative obscurity of otherwise public events of these sorts has been found to have a unique status in domestic U.S. law, even while specific kinds, such as telephone call logs (or "luds") were deemed to be in the public domain, with no presupposition of privacy attached, in a 1979 Supreme Court ruling.

[6] This surveillance, and its impact on both the surveilled and the surveillance personnel, was detailed in the moving and troubling film, "The Lives of Others" (*Das Leben der Anderen*, 2006).

[7] The relative tolerance of the use of deadly force for purposes labeled as "self-defense" is criticized on a number of grounds by David Rodin, in *War and Self-Defense* (Oxford: Oxford University Press, 2003), and in Jeff McMahan, *Killing in War* (Oxford: Oxford University Press, 2009).

[8] As Neil Rowe, a computer scientist and cyber expert, first pointed out in "War Crimes from Cyberweapons," *Journal of Information Warfare* 6, no. 3 (2007), pp. 15–25.

[9] Daniel J. Solove, "The End of Privacy," *Scientific American* (September 2008), pp. 101–106, libserver.wlsh.tyc.edu.tw/sa/pdf.file/en/e080/e080p110.pdf.

[10] In a star-crossed and ill-timed *Amsterdam Law Forum* article—published, ironically, on June 5, 2013, the same day that the first of Snowden's revelations appeared under reporter Glenn Greenwald's byline in the *Guardian*—I argued simply that the security of rank and file citizens of all nations might require some relaxation or compromise of the norm of privacy in the cyber domain, not realizing that this conflict had been settled in favor of greater security by agents of our government, out of sight, and thus lacking the full knowledge and consent of those allegedly being protected. See "Privacy, Anonymity, and Cyber Security," *Amsterdam Law Forum* 5, no. 2 (2013), pp. 107–14, ojs.ubvu.vu.nl/alf/article/view/311/485.

[11] David E. Sanger, "Differences on Cybertheft Complicate China Talks," *New York Times*, July 10, 2013, www.nytimes.com/2013/07/11/world/asia/differences-on-cybertheft-complicate-china-talks.html.