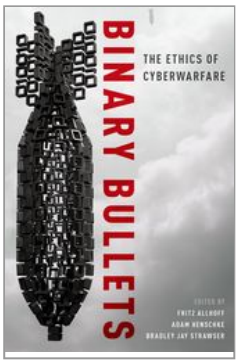

Binary Bullets: The Ethics of Cyberwarfare
Fritz Allhoff, Adam Henschke, and Bradley Jay Strawser

Print publication date: 2016
Print ISBN-13: 9780190221072
Published to Oxford Scholarship Online: December 2015
DOI: 10.1093/acprof:oso/9780190221072.001.0001

Moral Concerns with Cyberespionage

Automated Keyword Searches and Data Mining

Michael Skerker

DOI:10.1093/acprof:oso/9780190221072.003.0013

Abstract and Keywords

This chapter addresses the morality of two types of national security electronic surveillance (SIGINT) programs: the analysis of communication "metadata" and dragnet searches for keywords in electronic communication. The chapter develops a standard for assessing coercive government action based on respect for the autonomy of inhabitants of liberal states and argues that both types of SIGINT can potentially meet this standard. That said, the collection of metadata creates opportunities for abuse of power, and so judgments about the trustworthiness and competence of the government engaging in the collection must be made in order to decide whether metadata collection is justified in a particular state. Further, the moral standard proposed has a reflexive element justifying adversary states' intelligence collection against one another. Therefore, high-tech forms of SIGINT can only be justified at the cost of justifying cruder versions of signals intelligence collection practiced by a state's less-advanced adversaries.

Keywords: metadata, privacy, surveillance, reciprocity, automated searches, cyberespionage

States engage in espionage, including cyberespionage, as part of a continuum of actions to pursue their national security. This chapter will address the moral permissibility of two types of remarkable electronic intelligence collection that former National Security Agency (NSA) contractor Edward Snowden charged the NSA and the Government Communications Headquarters (GCHQ) with undertaking: keyword searches, in which automated collectors record electronic communications anywhere in a targeted region containing phrases, words, or names of intelligence interest; and metadata analysis, in which the pattern of communications in a particular region is mapped. I develop a standard for assessing coercive government action based on respect for the autonomy of inhabitants of liberal states and argue that both types of signals intercepts (SIGINT) described by Snowden can potentially meet this standard. That said, the collection of metadata creates the conditions for some unsavory government behavior and so judgments about the trustworthiness and competence of the government engaging in the collection, as well as the threat level the state faces, must be made in order to decide whether metadata collection is justified in a particular state. Further, the moral standard I propose has a reflexive element justifying adversary states' intelligence collection against one another. Therefore, high-tech forms of SIGINT can only be justified at the cost of justifying cruder versions of signals intelligence collection practiced by some technologically advanced states' less-advanced adversaries.

12.1 Politically Legitimate Forms of Coercive State Action

The novel methods of espionage made possible by cybertechnology resist easy comparison with more traditional methods of spying. We therefore need to be clear **(p.252)** about the moral foundations of espionage and military operations (discussed in this section) as well as the moral right(s) potentially violated by keyword searches and data mining (see section 12. 2). Section 12.3 will distinguish two morally different types of information gathering helpful to refine our argument about the SIGINT techniques discussed in section 12.4.

I will stipulate certain starting assumptions relevant to a moral foundation for espionage.[1] In doing so, I will be assuming a liberal political system as the background for the relevant domestic and international intelligence operations. I will also inquire as to what sort of intelligence operations are "politically legitimate" in these states—that is, what government actions are just, given the government's liberal underpinnings.[2] Politically legitimate actions do not in principle violate the rights of inhabitants of the state even if these actions are coercive in nature (I use this term broadly to refer to actions or laws that compel people to do things they do not otherwise want to do, e.g., tax collection, business regulation, arrest, prosecution, etc.). The criterion for politically legitimate state actions is consent-worthiness by the inhabitants of that state.[3] To judge whether some policy is consent-worthy, the theorist first conceives of an abstract consenter (with a particular moral constitution), then judges whether consent to the policy is logically necessary to protect the consenter's moral constitution; or, to put it another way, whether it would be self-contradictory to dissent to the policy given that the policy contributes to the consenter's protection. Since we are concerned here with national security actions, and thus chiefly with securing the negative freedom of a state's inhabitants (i.e., freedom from rights violations), we can operate with a fairly simple version of autonomy characterizing the abstract consenter. Since many versions of autonomy assume that negative freedom is a precondition for any more complex expressions of autonomy, my theory about politically legitimate intelligence operations should be insertable into political theories using more complex and detailed models of autonomy than I will develop here.

The model I use sees the consenter's autonomy expressed in specific arenas of thought, speech, and action in the form of rights. All abstract consenters are considered morally equal. I reject an atomistic (e.g., Hobbesian) model of autonomy that sees people as naturally autonomous outside some kind of settled political **(p.253)** community. Such a model conceives of government coercion as existing in tension with citizens' natural autonomy, a tension that is tolerated in exchange for the conveniences of living in a state.[4] Rather, the model I use sees autonomy as a nested concept, entailing as a necessary background an environment relatively free of rights violations and, so, a law-governed political entity (a formal state, in most empirical instances) with the coercive power to prevent and punish rights violations. Such a political entity is a necessary material precondition for a group of people to enjoy the full realization and expression of their rights over time consistent with equal rights expression. This is because an environment free of intentional or inadvertent rights infringements is a precondition for the realization and expression of one's autonomy in a given moment and over time. For example, one cannot build a house if one is being attacked or robbed of one's tools; and one would not even plan to build a house if one could not trust one's neighbors not to destroy or occupy it. Further, one will not develop the psychological faculties necessary for positive freedom (the capacity to deliberate on one's own and craft plans) if one is constantly in fear and want. So, provided certain constraints discussed below, coercive government actions, such as police activities, are in keeping with inhabitants' autonomy even when they restrict a person's autonomy in a particular instance. This follows, because the government actions are aimed in total at creating the environment relatively free of rights violations necessary for inhabitants to enjoy the full realization and expression of their rights consistent with universal and equal enjoyment.[5]

The underlying purpose of protecting inhabitants' autonomy creates the grounds for rejecting both certain government actions that are very harmful to autonomy and strategies meant to create an environment *perfectly* free of rights violations (because such strategies will likely cause intolerably high levels of rights infringements). The preferred moral framework I call the "security standard" endorses government tactics surviving a three0stage winnowing process. The standard (1) canvasses locally available tactics aimed at securing an environment relatively free of rights violations or the threat thereof; (2) isolates the most reliable, efficacious, proportional, and efficient tactics of those locally available; and (3) endorses the most rights-respecting among the tactics meeting the practical metrics of (2).

We can assume that any autonomous person would consent to domestic government actions aimed at securing a domestic environment relatively free of rights violations meeting the security standard. This consent will also justify actions by military and intelligence operators aimed at creating a domestic environment **(p.254)** relatively free of rights violations by defeating external threats to a state's security. Since all people in the world can be modeled as consenting to a regime of outward-facing security-seeking actions, a model consenter's consent to foreign operations by her security services also potentially justifies action by foreign agents targeting her. This dynamic can best be explained by discussing its domestic parallel. Hypothetical consent is permissive when it comes to the justification of police tactics meant to keep the model consenter safe. Considerations of how to secure the safety of a model consenter justifies a series of actions aimed at rights violators or potential rights violators. At the same time, a principle of reciprocity, justifying police behavior targeting the consenter if the consenter is suspected of perpetrating or planning rights violations, urges that police exercise restraint. So the consent that we imagine autonomous people extending to domestic security-seeking tactics takes into account that they might be the target of those tactics. The same reflexivity must apply to outward-facing security-seeking tactics since the security standard references an abstract autonomous person rather than a person of a particular nationality. By hypothetically consenting to outward-facing actions directed at foreign security threats, one would give leave to other governments to engage in outward-facing actions directed at foreign security threats, including oneself, if one is reasonably perceived to be a security threat. One cannot be modeled as consenting to illiberal governments, perhaps led by paranoid or sectarian leaders, monitoring foreign citizens who do not plausibly present a national security threat. To be clear, this equitable treatment of foreigners is arrived at by a different consent-based route than equitable treatment of one's neighbors. Whereas the actions of one's domestic law enforcement agencies can potentially be justified when such actions contribute to the necessary conditions for autonomy in one's own state, an adversary foreign state agent usually is not working to maintain conditions of autonomy in one's state, but rather the opposite. We can see that it would not make sense to model hypothetical consent to foreign agents' work if we also claim to justify domestic agents' actions opposing these foreign agents. Therefore, adversary foreign state agents' actions are potentially justified indirectly, as an entailment of consenting to one's own agents' outward-facing actions. By way of analogy, if one hires a lawyer to sue someone, one cannot begrudge the target of one's lawsuit hiring a lawyer to defend her rights and interests.

This reflexivity should encourage a conservative attitude toward intelligence collection from particular suspected foreign targets. The model consenter must use herself as a reference point, asking whether she can consent to her state agents using tactics abroad that, via the principle of reciprocity, she must also permit foreign agents to use against her. Using this approach, the rule of thumb should be that security agencies should use the same information-gathering tactics abroad that they use domestically. For example, if the security standard indicates that warrants issued by judges are necessary for a security service to intercept a particular domestic inhabitant's communications, the same treatment should apply to a foreigner **(p.255)** targeted by the security service. Note, this standard marks a serious departure from current American practice, for example, where foreigners and US residents are subject to markedly different SIGINT practices and bodies of law.

That said, practical limitations on foreign agents acting abroad or the different nature of the target might suggest different tactics than their counterparts would use domestically, leading to greater or lesser infringements on the target's rights. For example, it might not be as feasible to have ground units watch a suspected militant in the Swat Valley or the Ugandan rainforest as it would be in a domestic suburb. This limitation might prompt the surveilling agency to use airborne surveillance platforms, which might be more privacy-infringing than ground-based options in that they can see over walls and into compounds ground units cannot. To say this more privacy-infringing tactic is consent-worthy under the security standard is to say the model consenter permits her adversary's security agencies to attempt the same in her country if it confronts the same practical limitations there.[6] I will return to this point later in the chapter.

A wide range of concrete practices could be justified if the security standard permits security services to conduct foreign operations employing the most reliable, efficient, rights-respecting, and so on tactics available to that service. The best locally available tactics justified by the security standard will vary depending on a given state's wealth, size, technological prowess, and ingenuity. If the standard then effectively permits all security actors to "do their best," the standard allows situations in which, for example, wealthy country A's intelligence services can conduct very discriminate, sophisticated, targeted, and automated intercepts of foreign intelligence targets' communications—so that very few innocent people have their privacy infringed or violated—while also permitting poor adversary country B's intelligence services to conduct relatively crude, indiscriminate intercepts that infringe on the privacy of far more innocent people. As an example of crude intelligence gathering, an American NGO employee, previously posted in Uzbekistan, told me that the Uzbek National Security Service (NSS) listens to and tapes all visiting foreigners' phone calls as a matter of course. Yet the NSS only has the capacity to record thirty minutes at a time on its antiquated analog equipment and so simply disconnects calls on the thirty-first minute.

Intelligence collection activities fail the security standard in particular instances if one state's adversary's best methods of intelligence collection are so crude as to be imagined to be intolerable to the inhabitants of the target state. In this case, intelligence officers would need to refrain from collecting from a certain state if their behavior would justify retaliation by the target state engaging in its crude collection methods. By way of analogy, military actions against a **(p.256)** state fail the security standard even if otherwise just when the target state's only method of defense is use of WMDs.[7] That said, unlike military cases, it is difficult to think of an example of SIGINT that would be so rights-infringing as to be intolerable for any state to tolerate at the hands of its dangerous adversary if that was the price of garnering signals intelligence. Crude forms of SIGINT might not be consent-worthy if the reward for the risk was lower, such as if the target state did not pose a military threat to the collector state. The Uzbek case falls between these two clear extremes. Western states do have some security concerns in Uzbekistan potentially warranting intelligence operations directed at state and nonstate actors; the NSS apparently does not have the resources to monitor the communications of citizens of western states; and there are few western expatriates in Uzbekistan. All told, the security standard can probably justify western SIGINT operations against Uzbek targets. So watch out next time you are in Tashkent.[8]

12.2 The Right to Privacy

The signals intercept operations and accompanying analysis under discussion in this chapter are not as destructive as traditional military actions. Yet they are deeply troubling for their presumed infringements on people's privacy. Before discussing the tactics in detail, we need to clarify what is meant by infringements on, and violations of, privacy.

There are two definitions of mental privacy commonly used by philosophers: (1) a mental space of one's own, safe from external intrusion or disruption; and (2) a power to control the revelation of personal information. A certain mental **(p. 257)** space of one's own is thought to be a precondition for moral autonomy.[9] "Privacy is the condition of having secured personal space, personal space is space required to reason, and individuals have a fundamental moral right to reason as a means of securing autonomy."[10] We would not be able to plan for the future, develop a sense of self, or control our interactions with others if every thought was exposed prior to our decision to share it publicly.[11]

The power to control personal information helps protect the mental space where one should be free to reason and reflect. One would alter one's behavior, conversations, reading habits, and thoughts if one was concerned that one was under surveillance and, thus, was being forced to reveal ideas prior to their maturation.[12] The power to control personal information also puts one in control of one's intimate relationships, which are made intimate, in part, by the decision to divulge personal information to certain people.

Certain information is kept private because knowledge of it gives the knower power over the target. The information need not be what many cultures consider inherently private, such as information regarding the body, health, money, and sexuality, but also aggregated mundane information that in sum gives a portrait of the target's daily life.[13] Some private information can be damaging to the target in specific ways on account of the structure of society—leading her to lose her job, marriage, security clearance, health insurance, the trust of others, and so forth—and some information could be damaging if particular people wanted to harm her. For example, someone who wants to attack the target or rob her house would find her daily schedule of special interest. More broadly, the agent's collecting private information erodes the target's autonomy. There is now an asymmetry of knowledge and power between the target and the agent. He knows information about her normally only revealed to a friend, relative, or lover, but he is not any of those things. She did not choose to reveal this information to the agent, even though this is information she ordinarily only chooses to reveal to intimates. This knowledge can give the agent leverage over her in many interactions, as he **(p.258)** can use that knowledge she does not know he has to shape her perception of him, manipulate her, and make her irresistible offers. The asymmetry of knowledge is problematic even if it is not leveraged into some kind of invidious action. The agent has effectively coerced a level of intimacy from the target and taken from her the opportunity to choose how to present herself to him.[14] Thus, she is wronged even if she does not know about the privacy intrusion and even if the agent does not use the information to directly harm her.

In defining a privacy violation, I will focus on the second definition of mental privacy discussed above, the power to control personal information. I focus on this definition obviously because various kinds of communication intercepts remove the power to control personal information from the intelligence target. The first type of privacy (mental space) violation may also occur if the surveillance is discovered or assumed by the target since the knowledge that she is being watched will burden her thoughts.[15]

Having clarified the moral importance of privacy, we can now discuss the difference between harms, infringements, and violations related to privacy. Distinguishing these three things will be important to identify what is problematic about the cyberespionage techniques under discussion. One can suffer a harm in the arena covered by a right like privacy without being the victim of an infringement or violation. A woman suffers a harm associated with the involuntary disclosure of private information if she drops her purse in front of a coworker and some sensitive items spill out. The purse-owner feels the embarrassment, and the coworker has the knowledge that is normally associated with a rights infringement or violation. However, infringements definitionally involve an external imposition of harm or limitation through another person's action. Infringements that are not accidental, excused, or justified are rights violations, breaches of the agents' duties to respect others' rights.

I suggest a three-point definition of privacy infringement. First, privacy infringement involves the collection of a significant amount of information about the target that the target would not ordinarily reveal to a stranger. The context, the intent of the agent, and the choices of the target regarding what she considers private determines what is a significant amount. As discussion will clarify below, one datum might suffice if it is the sort of thing that most cultures consider private, such as information pertaining to the body, health, sexuality, and wealth. In other cases, aggregated facts—each innocuous on its own—together may reveal a portrait of the target she would not share, *as a whole*, with strangers or with certain strangers.[16] To be clear, **(p.259)** while the aggregation of many mundane details might amount to an infringement, the collection or observation of a single detail would not.

Second, in order for the privacy infringement to be actual rather than potential, the information must be attached to a particular person. A person gains no power over anyone if he finds an unaddressed love letter on the street. He knows someone's heart is aflutter, but cannot use this information to gain any advantage over or insight into any particular person. Third, the sensitive information has to be eventually known by a person. This is in part a reflection of the normative fact that rights infringements are actions of people. Falling boulders, machines, or wild animals can harm people, but not infringe on their rights, since boulders, machines, and nonrational creatures are not capable of self-limitation of their actions based on a rational appreciation of a mutual web of rights and duties. A machine might scan and record someone's sensitive information, but the machine's storage of this information does not create an asymmetry of power between itself and the target, again, because the machine is not implicated in the web of rights apportioned equally by theory to each adult person. The elements of an infringement may be present when we consider the agenda of the human designer of the machine and the human analyst who may study the stored information.

12.3 Surveillance and Patrol

It is useful to introduce a distinction between two actions, patrol and surveillance, in order to understand the moral significance of keyword searches and data mining. It will be helpful to use examples of patrol and surveillance involving human observers as models for cyberoperations since we have stronger intuitions involving the former. In patrol, the agent monitors a particular area, alert for suspicious behavior or other types of danger. Patrol is not focused on a particular person. Examples of patrol might include a policeman walking a beat, an air marshal sitting on an international flight, or a naval task force sailing back and forth in a commercial shipping lane.

Patrol raises far fewer moral concerns than surveillance when conducted by a just liberal state. First, a state agent's patrolling is merely a more concerted expression of what everyone does every day: observing activities occurring in public view in one's immediate vicinity and reserving the right to respond if something untoward or dangerous appears to be happening. The state agent has a special obligation to do what the ordinary person has a permission and weaker duty to do in the event of some emergency. Second, the patroller's attention is not focused on any particular person, and so patrol does not trigger the moral concerns related to infringements and violations of a particular person's privacy. The patroller is not gaining power over a particular person; he is not taking a prurient interest in a particular person; nor is he doing anything to make a reasonable person feel threatened (on the contrary, the **(p.260)** presence of law enforcement officers might well comfort a person in a just state). Third, the observed parties do not have a right not to be observed by a patrolling agent. The patrolling agent in a just state is not doing anything untoward. By exiting their homes, the observed parties tacitly consent to being observed by people on the street.[17] The patrolling agent is visible—often identified as a state agent—and so the observed party tacitly consents to be seen by him in particular. They cannot be modeled as tacitly consenting to be observed by an undercover state agent per se, but again, the patrolling agent is engaging in a permissible activity. It is reasonable to postulate an expectation of "privacy in public"[18]—a desire to not be closely observed in activities we perform in public but wish to keep private, such as shopping at a pharmacy or reading a letter on a train—but in most cases, the patrolling agent would not trespass this ambiguous border of privacy. While he might be more attentive than the layperson and this attention might press the boundary of privacy in public (e.g., a policeman might scrutinize the unusual bulge in someone's clothing or an airport guard might deliberately look at each face she sees for a second), it can usually be justified as a protective action in a just state. All these comments are restricted to patrols in a just liberal state. In a tyranny or other type of unjust state, political power is used to oppress inhabitants or a portion of the population for the benefit of the ruling clique or a privileged group. In this context, mere patrol serves to remind inhabitants of the scope of the government's power.

In contrast to patrol, surveillance raises a host of privacy concerns. Trailing someone, intercepting her communications, and watching her in and outside her home provide the agent with a profound degree of knowledge most cultures consider private. The agent gathers two kinds of information that strangers ordinarily do not know about one another and that people do not ordinarily reveal to strangers. First, the agent learns things that go on in the privacy of the target's home and in her communications—both occasions when she assumes her actions and words are private, only revealed to those she chooses. Second, the agent aggregates public actions to give a full picture of the target's daily life that no stranger (who might see her in a given moment) would know. While the individual elements of the target's daily public schedule are not necessarily sensitive (e.g., she picked up her dry cleaning, she bought coffee), their aggregation as "her daily schedule" is sensitive because it can give the agent significant power over the target.

Surveillance is a graver matter than many discrete privacy violations because it is definitionally expressive of a broader intention on the agent's part than intentions associated with discrete violations. It is difficult to think of benign reasons for ordinary citizens to engage in surveillance. Whereas a discrete privacy violation may result from the agent wanting to see the target naked, or find out a specific piece of **(p.261)** information about the target, the purpose of surveillance is to develop a portrait of the person, potentially inclusive of every facet of her life. This account of the target's life amounts to a major privacy violation, because the agent has gained the power that comes through unilateral knowledge of personal information with respect to (nearly) every facet of her life, not just one facet of her life (say, regarding her commercial habits).[19] In fact, whereas discussions of privacy are complicated by the fact that the boundaries of what is private are culturally constructed, surveillance would appear to be problematic in most cultures because it observes so many activities that might be considered private and because it attains knowledge of the target's life profile.

12.4 Tactics

We will now consider two types of intelligence-gathering operations potentially infringing on or violating their targets' privacy. I will speak about these operations at a certain level of generality, without ascribing them to specific agencies. This, because post-9/11 reports on the activities of various intelligence agencies are inconsistent, fragmentary, and frequently disavowed by people who are both in the position to know the truth about the operations and incentivized to lie about them.[20] Historically, initial reports of clandestine government operations often prove to be inaccurate. So we will consider two tactics as ideal types, with a presumed family resemblance to actual operation, past, present, or potential.

Keyword Searches—There are automated programs in existence that scour communication networks, collecting communications transmitted through fiber-optic cables or the electromagnetic spectrum. Supercomputers scan the intercepted data for “selectors”: certain words, names, or phrases associated with potential intelligence targets. This form of data mining is different than traditional wiretaps, pen registers, and pen traps and traces in that these keyword searches are not directed at particular suspects and do not necessarily require the participation of the phone company to physically manipulate the routing of calls. Rather, keyword collections are more like vacuum cleaners that collect everything that is in the air and on the cables. Communications deemed of interest according to some automated algorithm are recorded and forwarded to human analysts who read them and decide whether they should be purged or forwarded for further intelligence analysis.

It will be helpful to separate consideration of keyword searches into automated search, analysis, and investigation phases. It seems important to segment the tactic **(p.262)** in this way—eventually referring to “keyword search and SIGINT-prompted investigation” to encompass the full action—since many people affected by the tactic may only be touched by the first or first two phases. Just the same, the program has to be assessed in its totality since the first two phases exist to collect data for exploitation in the third phase, and it is the third phase that is potentially the most controversial. Regarding the initial search now, I will consider whether this type of data mining amounts to a privacy infringement and, if so, whether this infringement can be justified.

The initial search does not meet the three criteria for a privacy rights infringement described in section 12.2. The initial capture of an email, blog post, cell phone conversation, or text with a flagged phrase in it captures what will often be a quite small amount of data, akin to a sentence one hears walking past someone who is talking on the phone. The whole communication is not yet read in a keyword search, but merely tagged and stored because of the suspect phrase. Second (and even if the intercept is a fairly comprehensive and self-contained communication), the communication is not attached to a particular person: it is merely associated with a phone number or ISP number. We can see that these first two elements do not necessarily amount to an infringement on privacy because the suspect communication may not even come from a human being; data of this sort could be an automated message sent by a computer. Third, no human has seen the communication yet; a computer scanned the communication and stored it. To the computer, of course, the suspect phrases are not even words, just electrons moving in a certain pattern. While this exposure might prompt someone in the target region to feel his rights have been infringed or violated since his private message was “opened” prior to its reception by the intended recipient, I think this is an emotional residue connected to the symptom, the harm, rather than the substance of a privacy right infringement or violation. High technology forces us to draw analogies based in the physical world; the closest analogue here, a spy opening one’s letter, diverges in too many ways from the SIGINT tactic under discussion to be helpful. On a closer view, the components of a privacy infringement are absent with keyword searches. A human being who can understand what she is reading is not opening a complete, identifiable piece of correspondence.

The reader may object that my focus is too fine-grained here. There may be a genuine objection to the government *trying* to infringe on one’s privacy by whatever means—the automated keyword search assembles some of the components of an infringement—and so we need to address below whether gathering the data making infringement possible is politically legitimate. A partial, preliminary response to this concern notes that to the extent that keyword searches are like patrols, the government is not trying to infringe on a particular person’s privacy. Keyword searches are like patrols in that they are passive forms of collection, only leading to more invasive activities when something apparently untoward is observed. The difference between keyword searches and traditional patrols is that the latter are restricted to observing public **(p.263)** activity, while the initial, automated keyword searches reach “into” messages the sender and receiver presumably intended to be private. As already discussed though, the “virtually invasive” nature of these collection methods alone do not necessarily amount to rights infringements. We will now consider the second stage of keyword search collection involving analysis of intercepted data by intelligence analysts.

Keyword search collection may provide enough information to amount to the first criterion of a rights infringement by the time a human being reads the intercepted communication. The analyst may listen to a conversation or read an email or text of sufficient length to give context to the suspect phrase and reveal sensitive information about the intelligence source. The automated program may also be programmed to collect a string of communications from the suspect source and so provide the analyst context in that way. In these events, the analyst still does not gain power over a specific person because the information is still likely associated only with a phone or ISP number (probably coded, at that) rather than a particular identified person. The analyst knows someone, somewhere, has been up to no good, but at this level, it is simply an account of actions without a connection to a specific person. From the analyst's perspective, the narrative would be indistinguishable from a dummy source—a copied passage from a spy novel or that unaddressed love letter found on the street—forwarded by his supervisor to test his analytical skills. Thus, keyword searches still do not involve rights infringements when the human analysis stage is included.

There will be a point with some collections, after a certain amount of aggregation, when the intercepted communications are disturbing enough to warrant further investigation into the source, utilizing all manner of investigative and intelligence collection techniques. At this point, other analysts and investigators likely become involved. Now all the criteria of a rights infringement are present: aggregated information is tied to a particular person and the information is read by human beings. To be clear, what constitutes an infringement is the investigation, which draws on, but goes beyond, the initial keyword search. We now need to consider possible justifications for this kind of rights infringement in order to determine whether SIGINT-prompted investigations are politically legitimate and so presumptively not rights violations.

Whether an action is a rights violation depends in part on the rights of the person affected by the action. The criminal or unprivileged irregular combatant[21] whose operational communications are intercepted does not have his moral rights violated because he lacks a right to contribute to criminal operations via those communications. Adversary military, privileged irregular combatant, or intelligence personnel have a right to discuss their operational plans with colleagues, since (according to **(p.264)** the traditional post-Westphalian just war tradition), these professionals do nothing legally or morally wrong in pursuing the national security goals of their states or nonstate entities. Yet since their adversaries have the same right to pursue the national security goals of their own states, those adversaries can engage in strategic behavior such as intercepting their enemy's communications.[22] The targeted service member or intelligence officer, therefore, is not wronged by having his operational communications intercepted. Further, assuming that the operationally significant information collected in the data-mining operation regards state secrets, foreign security personnel do not suffer personal privacy violations when their communications are intercepted any more than soldiers whose rifles are taken by the enemy suffer private property right violations.

Clearly, the cases of concern with keyword searches are the false positives, cases where the communications of innocent people are collected when their out of context remarks trigger automated collection. We have to focus on these false positives rather than legitimate intelligence "hits" if we are going to assess the political legitimacy of keyword searches, since even grossly inefficient and brutal tactics like arbitrary arrest and torture can occasionally stumble across a legitimate intelligence source. Rights-infringing investigations of suspects can be justified when they meet the security standard of being the practically best and least rights-infringing tactics locally available. A certain error rate is in principle permissible since security officials would not be doing their job if they only investigated known threats, to the exclusion of anticipating future threats. So intelligence agencies must likely, and may in principle, engage in some kind of collection from *suspected* intelligence sources in order to meet their mission of contributing to national security. The mandate to pursue suspected intelligence sources entails that some innocent people will be targeted.

So we need to consider if keyword searches and SIGINT-prompted investigations can meet the security standard of being the practically best and least offensive to targets' rights locally available. We will first consider how the tactic measures up in terms of deferring to targets' rights. Scholars lacking security clearance are somewhat hampered in considering this question because they are ignorant of other possible types of modern SIGINT. As already argued, the automatic collection and initial human analysis phases of the tactic do not amount to rights infringements. Keyword searches and SIGINT-prompted investigations do compare favorably in terms of rights deference compared with older methods of collection, like steaming open envelopes or tapping phone lines, in that these methods identify specific people and utilize human analysts in the first instance. Keyword searches are also more rights-respecting than human intelligence (HUMINT) collection designed to accomplish the same goal of identifying suspicious communication the state might **(p.265)** not otherwise know to seek. HUMINT is morally fraught given that it often involves the corruption of the asset, the suborning of disloyalty, deception on the part of the recruiter, and great danger to both the asset and the recruiter. In the absence of clearly preferable alternatives and given the noninfringing nature of the first two phases of the tactic, I will tentatively say that keyword searches and associated investigations are sufficiently consent-worthy to meet the rights element of the security standard.

Having addressed the rights-respecting aspect of the security standard, it remains to be considered whether keyword searches meet the security standard's practical aspects of efficacy, reliability, efficiency, and proportionality. Regarding the program's efficacy, it has to be considered that imagery analysis—the analysis of imagery collected by satellites or reconnaissance aircraft—cannot substitute for signals intercepts since communications about military or terrorist operations are not necessarily accompanied by simultaneous physical actions. HUMINT can provide the same kind of behind-doors information about targets' communication, but is less efficacious than signals intelligence in several respects. HUMINT can be expected to be resource-limited compared to SIGINT since developing human sources is labor- and time-intensive and not consistently fruitful. It may be extremely difficult to cultivate human intelligence assets in all the locations one desires because certain government programs or installations employ small numbers of dedicated, highly vetted, highly monitored people intelligence officers will have great difficulty locating, much less "turning." Some states or groups are so isolated from the outside world that penetration by intelligence officers is all but impossible. By contrast, any of these selective, secretive, isolated groups, installations, or states are potentially vulnerable to signals intelligence collection.

Further, even for highly competent, well-funded agencies, the scope of HUMINT operations is limited by prior intelligence collection. Intelligence agencies only know to try to cultivate or collect from assets associated with adversary organizations or installations with which they are already familiar. The comparative benefit of wide-scale SIGINT is that it can alert agencies to threats they did not know were germinating.

The reliability of keyword searches is difficult to assess without disclosure of the types of searches conducted and the standard yield of useful intelligence they produce. Elsewhere, I argue that specific information about SIGINT capabilities and search terms should be classified lest whole avenues of intelligence collection be shut down by adversaries.[23] Therefore, I am forced to compare keyword searches conducted by intelligence agencies with those performed by civilians using Google and the like. Assuming that there are actually a tiny number of intelligence targets whispering into their satellite phones about nefarious things relative to the total human population, the comparison with civilian search engines (e.g., imagine searching for someone with a common name) suggests that intelligence searches would yield relevant intelligence along with a huge volume of false positives. Given the amount **(p.266)** of time and manpower that would be necessary to analyze every initial intercept containing a suspect word or phrase, one imagines that further automated filtering occurs prior to a human analyst seeing any collected information, including longitudinal collection of other intercepts from the same source and cross-referencing with geographic areas of interests. Thus, it is reasonable to assume that an exponentially smaller number of false positives reach the desks of analysts than are generated in the original collection. Finally, it would seem that this filtering coupled with human analysis is a reasonably reliable method of selecting targets for more focused investigation. While human analysts no doubt miss the significance of certain messages in which targets speak in unfamiliar codes, or think some innocuous conversation is laden with code words, it is hard to imagine a more reliable method of analyzing communications than to have trained analysts read them.[24]

There are also difficulties assessing the efficiency of keyword searches without access to the classified details of such programs. While it is not possible for someone without access to the relevant classified information to know if there are more efficient contemporary techniques available, it is possible to make some speculative comparisons between keyword searches and known historical alternatives. It must be more efficient for supercomputers to analyze data in milliseconds than have corps of analysts steaming open envelopes or listening to phone calls in real-time. It is also reasonable to assume that the keyword searches conducted by the best-funded intelligence agencies are of the most or nearly most efficient types currently available because efficiency (the rate at which collected intercepts can be analyzed and flagged for security-sensitive information) is something readily measurable by engineers employed by the agencies and by the inspectors who oversee the engineers' work. The gap between collected and analyzed information is a knowable quantity and, presumably, a matter of concern to intelligence analysts. Given presumed institutional interests in ever-increasing efficiency and incentive structures for engineers pegged to customer demands, it can be assumed that well-funded agencies would be able to get the most efficient collection methods.

It is no surprise that the values relevant to a proportionality calculation are also hazy, though not quite to the degree as to reliability and efficacy. In this case, proportionality has to be assessed in two stages. The potential good done by the program has to be considered by assessing the evils avoided, in other words, the scope of the **(p.267)** threats posed by the state's adversaries. Then, the efficacy of the proposed program has to be assessed in order to know what percentage of the maximum potential good done can be theoretically accomplished. With respect to the first stage, applying the security standard in the case of intelligence collection is harder than applying it to preparations to meet a concrete threat, such as the threat of muggers in a particular neighborhood, because the "good done" element of the proportionality concern is undefined. It is difficult to know if one's security establishment is overdoing intelligence collection without knowing about the threats that are currently germinating. Yet this knowledge is only ascertainable through intelligence collection. That said, some crude estimates are possible: small, resource-poor countries with minimal international concerns (shipping, foreign bases, etc.) likely face fewer national security threats than large, wealthy, internationally involved countries. Countries that have been at peace for decades presumably face fewer threats than those engaging in antagonistic international actions or those hearing sustained, plausible threats from states or nonstate groups. Surinam or Bhutan, for example, probably face far fewer threats than say, Iran or Israel, and so are less able to justify a significant SIGINT collection capability. Still, due to the need to anticipate future threats, the security standard will justify collection efforts somewhat disproportionate to current, known threat levels.

Regarding the second stage of a proportionality assessment, selective disclosures by intelligence agencies as well as leaks present an ambiguous picture of efficacy of keyword searches at forestalling terrorist attacks or frustrating enemy military maneuvers. The public does not know about every counterterrorist or other type of military operation and does not know how many are predicated on key signals intercepts. And the public does not know how many intercepts relevant to security concerns were incorrectly analyzed or analyzed too late. No one knows the ratio of intercepted communications to the total number of security-sensitive communications sent between actors. Even if the “good done” portion of the proportionality calculation was clearly known, the calculation would still involve ambiguity because it involved a comparison of different values, say thousands of people’s communications intercepted and read every year compared to a few thousand lives saved per year on account of frustrated terrorist plots or disrupted military maneuvers. Still, since keyword searches are minimally invasive and non-rights-infringing forms of patrol, I am inclined to think that in a state facing significant national security concerns, the proportionality calculation would favor the prospect of saving human lives even if the number of communications intercepted increased exponentially. I suspect the same holds true for SIGINT-prompted investigations amounting to justified rights infringements in just liberal states—when the process of filtering and analysis determining whether a SIGINT target is disclosed to human analysts meets the security standard itself. In other words, SIGINT-prompted investigations would fail the proportionality test if the filtering process was so lax that huge numbers of innocent people were subjected to detailed surveillance and investigation. Keyword **(p.268)** search programs would also fail the proportionality element in low-threat environments if there was no way to ensure good behavior on the part of analysts—since analysts collecting or reading intercepts for puerile reasons would violate targets’ rights.

In sum, keyword searches and SIGINT investigations are designed to gather information other imagery collection and HUMINT cannot; the programs appear to be more reliable than conceivable wide-scale collection alternatives; it is reasonable to assume that the relevant programs conducted by the few states with the resources to engage in them are fairly efficient; and the programs can be proportional when professionally run in states facing significant national security threats. I stressed above that many of my practical assessments are tentative given the secrecy surrounding the relevant programs and possible SIGINT alternatives. However, I have met my goal of establishing an abstract framework into which empirical details can be added. A particular program will fail to meet the security standard, for example, if it has a very high false positive rate or if its analysts are poorly trained; if the state has little realistic need for such a program;[25] and when practically better and more rights-respecting programs become available.

Metadata Analysis—There are many different types of metadata analysis, performed by entities ranging from commercial firms to political campaigns. Certain intelligence agencies have recently acknowledged gathering “telephony” metadata from telecom companies. This type of intelligence does not reveal the content of communications but the time, duration, and phone numbers or ISP numbers involved. One possible use would involve mapping the social network of a suspect drawing on stored tranches of metadata seized for the entire population in a region and gathered years prior to when anyone in the tranche was identified as suspicious. In what follows, I will first address concerns regarding collection and storage of metadata generally, and then focus on particular concerns with the government collecting and storing it.

A metadata map of all the communications in a particular region is not greatly sensitive unto itself. Modern communication technology means that communication patterns are invisible to the naked eye. It does not follow that individuals enjoy a legitimate expectation to avoid appearing to others as part of a telecommunication pattern. It will be helpful to draw an analogy with visible communication patterns to defend this claim.

Imagine a rookie policeman manning his post at the main intersection of a small town on the Canadian-US border in the early twentieth century. From his post, he observes the pattern of American and Canadian neighbors going from house to house to visit one another. He notes their pattern of communication including **(p.269)** the origin and destination of communicators without knowing the content of their communication or the identity of the communicators (being new to the job and the town, he just notes general information like gender, age, height, hair color, etc.). For their part, there is no reasonable expectation on the part of the citizens not to be seen by others when they leave their houses in order to communicate with their neighbors. For the policeman's part, patrol by authorities is permissible, as discussed above, and a policeman's noting of a *pattern* in human traffic is merely an extension of patrol, in which discrete moments of the patrol day are linked together in the policeman's memory. The situation does not change if a machine employed by the police notes the pattern of communication. A machine's "memory," and perhaps its observation, will be vastly better than the patrolman's, but the citizens are still not specifically identified as individuals in such a pattern (machine memory will be discussed further below). Instead, they are just data points whose notable feature is taken from the pattern-noter's agenda. An observed person appears in the pattern as a "caller," "redhead," "pedestrian," "motorist," and the like. One does have a claim to control mundane pieces of information about oneself, but the pattern-relevant pieces of information are not being traced back to a unique person at this stage. So at this stage, so limited, the observation does not infringe on a person's privacy. The pedestrians have no right to demand that they not be objectified, in a sense, as a data point (e.g., as a pedestrian, voter, redhead, etc.) for someone else's observation or research. Again, it is worth emphasizing the distinction between patrol and surveillance. While there is no reasonable expectation against being noted as part of a pattern, there is a reasonable expectation against being surveilled (e.g., being followed all day), and in the course of this action, one's itinerary and associations being catalogued by a spy.[26] If being noted as part of a pattern is not rights infringing, we will address concerns related to the ways in which a machine-stored pattern could be later used below.

One cannot expect to go unnoticed when one visits another person's house in a populated area. So one arguing for the great sensitivity of telephony metadata needs to argue that patterns of communication become especially sensitive when the pattern is a product of technology permitting people to communicate without traveling from door to door. One needs to argue that the expectation of anonymity associated with such technology amounts to an expectation of privacy.[27] To be clear, one can be accustomed to anonymity, because of the nature of the technology, without that fact creating a normative expectation of privacy. There is an argument to be made that some people may reasonably expect privacy regarding the destination of at least their local calls when they choose to call instead of leave their homes to visit their **(p.270)** interlocutor. One can certainly imagine instances when a caller does not want even the destination phone number revealed to certain parties. The phone number of a man's mistress (recognizable to the philanderer's wife) or the widely advertised number for a phone sex company, abortion clinic, or local welfare office could be sensitive in this manner. Yet first, obviously, these are exceptional cases; and second, they are moot when the person does not have a realistic option of physically visiting the recipient of their communications. If the person has no alternative but to write, call, or email, we cannot assume that the sender wants to keep the recipient's address secret merely by virtue of the fact that he uses technology making his communications invisible to the naked eye.

We still have to substantively determine if telephony metadata is so sensitive that we can expect its collection would amount to a privacy infringement. I will argue that is not sensitive to this degree. One should understand that telecom companies have access to communication metadata, and one presumably endorses their automated monitoring of Internet and phone traffic to prevent the overloading of servers or other technical glitches.[28] True, people rely on telephones and the Internet, and so may only grudgingly tolerate telecom companies' possession of metadata. Yet the toleration of telecom companies having this information also suggests that metadata is not seen as all that sensitive. We can imagine that people would not use phones or the Internet, or would demand the immediate disposal of metadata, if use of communication technology or storage of metadata was seen as a violation of privacy on par with the public exposure of one's sexual habits or health information.

Yet one might object that a telecom company's possession of metadata is very different from a government's possession of metadata. We may understand that phone companies know the numbers we have been calling—just like we know our doctors have our health records—but also assume that phone companies do not use that information for invidious purposes.[29] While tacit consent legitimates the telecom companies' possession and analysis of metadata for the purpose of facilitating communication, the objection continues, intelligence agencies have a completely different interest in analyzing metadata.

Granting that exposing information to one party does not imply consent to universal disclosure,[30] an intelligence agency's gaining access to an arena without the target's explicit consent is not inherently problematic. The security standard, for example, justifies the state gaining access to normally private material such as one's communications, home, and possessions through a warrant process. The challenge in justifying broad metadata collection and storage by the government is that the target is not a suspicious party or clearly identified adversary state agent. Instead, the **(p.271)** target set is a broad swath of the population, collected against with the thought that someone in that group might be contacted by a foreign security threat in the future. Contrary to due process, evidence collection occurs prior to the identification of a suspect.[31] I will reply to this objection along with the "right to be forgotten" below.

To this point, we have seen that merely collecting metadata about physical movement (presumably associated with communication) is not rights infringing and that we cannot assume a different assessment when it comes to collecting telephony metadata. There is evidence that people do not see their telephony metadata as especially sensitive. We have yet to address due process concerns with a government collecting it. Before addressing those concerns, it will be helpful to entertain a critique of metadata collection focusing on the long-term retention of metadata as the problem rather than the inherent sensitivity of the information. This critique goes along with the due process critique of collection since data that is collected but not stored can hardly be used as evidence in contravention to due process.

Such retention, one might charge, trespasses against a reasonable expectation of the ephemerality of one's communications and associations.[32] If one has a right to define oneself to strangers, then one's actions, communication, and associations should be allowed to dissolve into the past, as it were. One would take from others the chance to define and present themselves if one kept a record of others' every statement and association, even if one did not consult the data in real-time but merely stored it for possible future analysis. Not only would retrospective analysis challenge a person's real-time self-definition, but the threat of one's interlocutors checking the data would likely inhibit one's present interactions, associations, and self-definition.

Yet there cannot be a generalized expectation against the retention of past behaviors and communication—an expectation that one's behavior dissolves into the ether—because one cannot make demands, practically or morally, on another person's memory.[33] Noticing patterns is essential to learning. Noting and remembering patterns of a person's behavior or speech might be essential to realize that a person is manipulative or untrustworthy. Further, observing patterns of past behavior or speech can lead a person to be deferential to a friend's sensitivities or to anticipate her needs.

Is there then a legitimate, *special* expectation against computerized capture and retention of past behaviors and communication? Critics of metadata collection by **(p.272)** intelligence agencies are not imposing demands on other people's memories but worried about digital storage. One arguing against metadata storage could point out we do not mind being passively observed by strangers in public because (a) there is an expectation that they will not ogle us and try to log every detail they see and (b) because we expect that they will not remember whatever they see for long. On one hand, people have a right to their memories; and on the other hand, those memories generally are not a threat to strangers. By contrast, automated data collection gets every pixel and syllable and can store that information for a near eternity.[34]

While this kind of record could be made for nefarious purposes, and could be abused even if it is made for good purposes, I do not think a near-permanent stored record poses an inherent normative problem. I will explain this point and then address the issue with potential abuse. Imagine an alien race interested in human beings for purely academic purposes. The aliens create machines that photograph and record humans' every moment. This information is studied merely to further the aliens' understanding of human beings and help alien academics make tenure. The aliens' creed forbids their interference with species on other planets. The aliens have sterling information assurance systems so that the collected information can never be leaked or used for purposes other than this kind of benign research. Neither privacy concern is in play here if the moral importance of controlling one's own personal information stems from the importance of choosing how to present oneself to others and protecting oneself from certain kinds of harm. The mere existence of a minute record of one's activities that will never be viewed by humans is not a problem. This leaves concerns about how the information will be used, and provided some justifiable use, how susceptible that benign process is to abuse. I will turn directly to these questions now and also address the due process question outlined above.

> The security standard is a framework for identifying politically legitimate government coercive actions—government actions meant to protect inhabitants through consent-worthy means. The standard can take into account the risk of abuse of processes in its proportionality calculation. The collection and storage of domestic metadata for the purpose of retrospective analysis of networks associated with particular foreign suspects appears to be a good candidate to pass the security standard. I will focus first on the rights-respecting element of the standard. First, there is no generalized right not to be seen as part of a pattern, as argued above, nor a generalized right not to be remembered as part of a pattern. Second, the collection of metadata does not trespass on very sensitive areas. Collection does not reveal the specific content of communications; particular people are not identified; and, presumably, a computer algorithm (rather than a person) collects and organizes the metadata. Third, the means of acquisition of the metadata is not burdensome to the public. No one's daily behavior is disturbed through metadata collection. **(p.273)**

> Storage and retrospective searches of metadata also meet the practical elements of the security standard. As with targeted searches, retrospective searches of stored metadata would likely score well in terms of reliability and efficacy. Storing metadata increases the efficacy of metadata searches because collector algorithms now have more data to sift, promising more opportunities to find contacts of foreign security threats. That said, the practice of storing terabytes of metadata with the idea that a foreign security threat might have used the relevant telecom networks sometime in the last few years is terribly inefficient, like storing all the hay in the world in case someone realizes he lost a needle in a bale a year or two back.[35] Yet there may not be more efficient methods. The numbers called on a particular line can be efficiently collected in real time with a pen register, but this collection does not help to recover the calls made in the past. The most efficient method for retrospective investigation is hardly reliable or even possible in many cases: that is, asking the suspect whom he contacted regarding a covert operation. Regarding proportionality, the good of being able to see whom a foreign security threat contacted domestically seems very important and the harm of data collection is slight unto itself. However, a significant concern on this point is not simply the collection for the purpose of possible future analysis regarding foreign security threats, but the possible use by a government for ordinary domestic law enforcement or political oppression. We will return to the question of proportionality below, taking into account this risk.

One can envision metadata analysis becoming as common a response to arresting a criminal suspect as the examination of fingerprint records or arrest records is now. Again, the concern is that unlike the collection of fingerprints, metadata collection would violate due process by collecting evidence prior to reasonable suspicion of a crime occurring.[36] Further, while fingerprints can reveal whether one was present at other crime scenes, metadata coupled with subsequent database searches can do more in providing a fairly full biographical sketch of a person. This critique stands even though the information collected is not very sensitive and the mode of collection is noninvasive, because it violates what some consider the proper relationship of the liberal state and its inhabitants. In the words of US Supreme Court Justice Arthur Goldberg, the government should "leave the individual alone until good cause is shown for disturbing him."[37] In order to maintain the general respect for the individual, the state should maintain such a distance that it has to labor when it comes time to prove guilt in the criminal justice arena. It should not be in a constant state of laying the groundwork for every inhabitants' future prosecution. Regarding concrete effects, knowledge that one's metadata could be used in future criminal **(p.274)** prosecution could have a chilling effect on one's innocent communications with people one suspects might be involved in criminal behavior and could certainly chill journalistic work or the sort of advocacy work that involves working with at-risk youth or ex-offenders. Some might harbor even graver fears that newly criminalized behavior could lead to retroactive prosecution via metadata analysis or that people could be persecuted for legal but politically unpopular views.[38] So a proportionality analysis has several significant risks to weigh against the possible counterterror or counterintelligence good done by metadata collection.

On these proportionality-relevant points, slippery slope arguments have to be tethered to a realistic assessment of risk, lest they be permitted to overturn every otherwise coherent argument and invalidate every policy prescription. I am sensitive to the argument that it is just plain unsettling that metadata might be stored by corporations or governments that can be used in combination with other archived information to develop detailed portraits of our lives. The concern is not necessarily over the inherent sensitivity of the information (all or most of it may be known by our relatives and friends) but the uncertainty of how it might be used by less trustworthy parties. In some states, the risks of metadata abuse may be high enough to make the possible harm both more likely and more costly than a terrorist attack or military or intelligence operation facilitated through contact with a domestic inhabitant. In other states, with strong histories of rule of law and cultures of responsible public service, the proportionality calculation may be resolved in the other direction. I do not think metadata collection and storage can be declared absolutely beyond the pale or absolutely essential to national security. Metadata collection offers its collectors a power promising certain benefits and bearing certain risks of abuse. Rather, I want to identify some of the possible risks and benefits associated with such a program and create a framework in the security standard for determining whether a metadata collection program should be implemented in a particular state given the balance of risks and benefits.

A practical solution addressing the due process and slippery slope arguments—and making us more confident of the proportionality of metadata collection and storage—would be for the government to refrain from collecting metadata and instead to require telecom companies to retain metadata for a set number of years.[39] The government could then obtain a warrant to collect the relevant metadata if it had a reasonable suspicion of a person. This maneuver appears to be in line with current due process procedures in some liberal states in which data more sensitive than metadata, like the content of calls and emails, can be obtained with a warrant. This solution no more admits of overreach by a government than any other **(p.275)** warrant process. Further, the proposed regime would only extend current practice, at least in the United States, since telecom companies are currently required to retain metadata for a period of time for regulatory purposes.

Conclusion

This chapter analyzed two types of SIGINT in reference to a moral framework based on respect for human autonomy. The security standard understands certain types of government coercion to support human autonomy by fostering an environment relatively free of rights violations. The standard assesses tactics by seeking a balance between achieving a positive security outcome and deferring to inhabitants' rights. The standard provides a framework for assessing whether the two types of SIGINT are justifiable in particular settings. The two forms of SIGINT may or may not be justifiable given the threats the state faces, the available technological alternatives, the professionalism of the analysts administering the program, and the justness of the government directing the program. Given the reflexivity built into the security standard, one can make an argument for the general, but not universal, permissibility of the two tactics. While a given tactic might pass the security standard on first analysis, the reflexive element of the standard permits adversaries to respond with a tactic that best approaches an in-kind response, even if the best they can muster is more rights-infringing than the tactics used by the more technologically advanced state. Thus, in particular concrete cases, an otherwise permissible intelligence-collection tactic may not be consent-worthy if the relevant adversaries' tactics are more rights-infringing than inhabitants in the first state could tolerate.

Bibliography

Bibliography references:

ACLU v. Clapper. S.D.N.Y. filed June 11, 2013. https://www.aclu.org/national-security/aclu-v-clapper-legal-documents.

Alfino, Mark, and G. Randolph Mayer. "Reconstructing the Right to Privacy." *Social Theory and Practice* 29 (January 2003): 1–10.

Apel, Karl-Otto. *From a Transcendental Semiotic Point of View*. New York: St. Martin's Press, 1999.

Bamford, James. *The Shadow Factory*. New York: Doubleday, 2008.

Benn, Stanley. "Privacy, Freedom, and Respect for Persons." *Nomos* 13 (1971): 1–26.

Bok, Sissela. *Secrets*. New York: Vintage, 1989.

Cohen, Jean L. "Equality, Difference, Public Representation." In *Democracy and Difference*, ed. Seyla Benhabib. Princeton: Princeton University Press, 1996.

Gross, Hyman. "Privacy and Autonomy." *Nomos* 13 (1971): 169–81.

Habermas, Jurgen. *A Theory of Communicative Action*. Boston: Beacon, 1985.

Inglis, John. “Remarks.” Presented November 22, 2013 at the Center for Ethics and the Rule of Law Conference, University of Pennsylvania Law School, Philadelphia, PA.

Kant, Immanuel. “On the Proverb: That May be True in Theory.” In *Perpetual Peace and Other Essays*, ed. Ted Humphrey, 61–92. Indianapolis: Hackett, 1983.

McCloskey, H. J. “Privacy and the Right to Privacy.” *Philosophy* 55 (1980): 17–38.

Murphy v. Waterfront Commission. 378 US 52, 55 (1964). **(p.276)**

Nagel, Thomas. *Equality and Partiality*. New York: Oxford University Press, 1991.

Nissenbaum, Helen. “Protecting Privacy in an Information Age: The Problem of Privacy in Public.” *Law and Philosophy* 17 (1998): 559–96.

Nozick, Robert. *Anarchy, State, and Utopia*. New York: Basic Books, 1977.

Rawls, John. *A Theory of Justice*. Cambridge, MA: Belknap, 1971.

Reiman, Jeffrey. “Privacy, Intimacy, and Personhood.” *Philosophy and Public Affairs* 6, no. 1 (1976): 26–44.

Scanlon, T. M. *What We Owe to Each Other*. Cambridge, MA: Belknap, 1998.

Simmel, Arnold. “Privacy Is Not an Isolated Freedom.” *Nomos* 13 (1971): 71–87.

Skerker, Michael. *An Ethics of Interrogation*. Chicago: University of Chicago Press, 2010.

Skerker, Michael. “A Moral Foundation for Government Secrecy.” Presented November 23, 2013 at the Center for Ethics and the Rule of Law Conference, University of Pennsylvania Law School, Philadelphia, PA.

Smith v. Maryland. 442 U.S. 735 (1979).

Solove, Daniel. *The Digital Person*. New York: New York University Press, 2006.

Stramel, James E. *Same Sex*. Lanham, MD: Rowman & Littlefield, 1999.

Waldron, Jeremy. “Theoretical Foundations of Liberalism.” *Philosophical Quarterly* 37, no. 147 (1987): 127–50.

Waldron, Jeremy. “Special Ties and Natural Duties.” *Philosophy and Public Affairs* 22, no. 1 (winter 1993): 3–30.

Westin, Alan F. *Privacy and Freedom*. New York: Atheneum, 1967.

Van Den Haag, Ernest. “Definition: The Nature of Privacy.” *Nomos* 13 (1971): 149–68.

Notes:

([1]) The ideas articulated in this section are developed in detail in my *Ethics of Interrogation*, chap. 2.

([2]) Certain intelligence operations by illiberal states can be justified to the extent that they are deployed for the same reason they are deployed in liberal states: the protection of the state’s inhabitants.

([3]) Inhabitants are the relevant consenters rather than citizens of states, because hypothetical consent is modeled in reference to abstract conceptions of the human person rather than in reference to people of particular nationalities. For examples of hypothetical consent, see Kant, “On the Proverb,” 61–92, 79; Rawls, *Theory of Justice*, 11; Waldron, “Theoretical Foundations of Liberalism,” 127–50, 138; Waldron, “Special Ties and Natural Duties,” 3–30, 25; Habermas, *Theory of Communicative Action*; Apel, *From a Transcendental Semiotic Point of View*.

([4]) E.g., Nozick, *Anarchy, State, and Utopia*, chaps. 2–5. See T. M. Scanlon, *What We Owe to Each Other*, 19–25.

([5]) To be clear, we might find a genuinely autonomous person outside a political community, marooned on a deserted island or living in a failed state like Somalia, but only if her formation occurred elsewhere. We do not expect children to grow up to be fully autonomous people in such environments.

([6]) The adversary agency’s permission does not mean agencies in the target state are not permitted to oppose their actions.

([7]) This argument might create perverse incentives for small states in particular to invest in WMDs at the cost of improving their conventional forces. Then larger states would leave them alone for risk of incurring an indiscriminate military response. Yet perhaps this incentive is not so perverse if it reduces the likelihood of war.

([8]) One might wonder if any states enjoy a unilateral right to collect against adversaries because of the illegitimate nature of the target government. Since the security standard is indexed to the protection of negative liberty, it justifies traditional policing and national security actions of even some illiberal and/or autocratic states. While the security standard does not justify repressive actions aimed at a government's nonviolent political or ideological opponents, it does justify the bread-and-butter responsibilities of a state aimed at protecting its inhabitants from street crime, piracy, terrorism, and foreign military attack. The security standard does not justify the coercive actions of states with governments that largely neglect ordinary inhabitants and use power largely to benefit ruling cliques. The coercive power of government is justified in order to create relatively crime-free environments for the benefit of the inhabitants of the state. As examples of states lacking the justification for coercive state actions, I would suggest: Amin-era Uganda, Mobutu-era Zaire, Duvalier-era Haiti, military-ruled Burma, present-day Equatorial Guinea, and North Korea.

([9]) Alfino and Mayer, "Reconstructing the Right to Privacy"; Stramel, *Same Sex*, 284, 285; McCloskey, "Privacy and the Right to Privacy"; Cohen, "Equality, Difference, Public Representation"; Nagel, *Equality and Partiality*, 142–43. Benn, "Privacy, Freedom, and Respect for Persons," 1–26, 3; Van Den Haag, "Definition: The Nature of Privacy," 149–168, 151. Alan F. Westin distinguishes four functions of privacy, one of which, "reserve," "protects the personality" by creating invisible walls between the person and the rest of the world. *Privacy and Freedom*, 32.

([10]) Alfino and Mayer, "Reconstructing the Right to Privacy," 10.

([11]) Bok, *Secrets*, 21–23. See also Westin, *Privacy and Freedom*, 34; Benn, "Privacy, Freedom, and Respect for Persons," 24–26; Simmel, "Privacy Is Not an Isolated Freedom," 71–87, 73.

([12]) Van Den Haag and Benn make similar points: "Definition: The Nature of Privacy," 151, and "Privacy, Freedom, and Respect for Persons," 10, respectively.

([13]) Nissenbaum, "Protecting Privacy in an Information Age," 559–96, 565; Solove, *Digital Person*, 146.

([14]) Gross, "Privacy and Autonomy," 169–81, 172; Van Den Haag, "Definition: The Nature of Privacy," 152.

([15]) Benn, "Privacy, Freedom, and Respect for Persons," 10.

([16]) Benn discusses related matters, "Privacy, Freedom, and Respect for Persons," 4–6; Nissenbaum, "Protecting Privacy in an Information Age," 565, 589; Solove, *Digital Person*, 43.

([17]) Reiman, "Privacy, Intimacy, and Personhood," 2–44, 44.

([18]) See Nissenbaum, “Protecting Privacy in an Information Age.”

([19]) See Solove, *Digital Person*, for concerns over the threat to privacy posed by the digital dossiers formed by the automated aggregation of mundane biographical details.

([20]) http://www.washingtonpost.com/nsa-secrets; http://www.theguardian.com/world/the-nsa-files; Bamford, *Shadow Factory*; Inglis, “Remarks.”

([21]) An irregular combatant is irregular in affiliation (belonging to a nonstate group) and/or tactics (using guerilla rather than conventional military tactics). An unprivileged irregular is one who fails the criteria for moral and lawful belligerency: obeying a unified chain of command, carrying one’s arms in the open, wearing identifying emblems, and obeying the laws and customs of war.

([22]) See Skerker, *Ethics of Interrogation*, chap. 7.

([23]) Skerker, “Moral Foundation for Government Secrecy.”

([24]) To compare the reliability of this form of SIGINT with one form of HUMINT, it strikes me that there would be fewer false positives with this sort of collection than produced in interrogations. The analyst at NSA or GCHQ or a similar agency likely has an overwhelming number of noninvidious intercepts come across his or her desk every day. Given the relatively “dumb” nature and global span of collection, the analyst presumably expects to get far more false-positives seized by computers than actual intelligence sources. By contrast, interrogators might well be prey to confirmation bias, particularly with suspected High Value Targets whom the interrogator reasonably assumes was only interdicted after thousands of man-hours of analysis, investigation, and tactical planning.

([25]) Programs that are weaker on the practical or rights-respecting fronts may be justifiable in states facing grave threats.

([26]) The ACLU appears to conflate knowledge about associations garnered through patrol and surveillance in their June 11, 2013 complaint in the Southern District of New York District Court. *ACLU v. Clapper*. https://www.aclu.org/national-security/aclu-v-clapper-legal-documents.

([27]) The distinction is George R. Lucas’s.

([28]) The US Supreme Court argued that people did not have a reasonable expectation of privacy regarding telephone metadata for this reason in *Smith v. Maryland*.

([29]) This argument exposes the shallowness of the Court’s reasoning in *Smith v. Maryland*.

([30]) Nissenbaum, “Protecting Privacy in an Information Age,” 585; Solove, *Digital Person*, 43.

([31]) One of the American Civil Liberties Union’s concerns is that absent a warrant, the collection of metadata by the NSA amounts to an unreasonable search. See legal documents relevant to *ACLU v. Clapper* at https://www.aclu.org/national-security/aclu-v-clapper-legal-documents.

([32]) The European Parliament passed a digital privacy law March 12, 2014 codifying a right to be forgotten. The law empowers citizens to force Internet companies to remove personal information from their servers.

([33]) See Nissenbaum, “Protecting Privacy in an Information Age,” 572.

([34]) Ibid., 576.

([35]) The analogy is one used by Inglis, “Remarks.”

([36]) It should be noted that some US states collect fingerprints of anyone who works with children, expressly for the purpose of tracking them if they are later accused of kidnapping or molestation. This practice would seem to offer fertile grounds of comparison with metadata storage.

([37]) *Murphy v. Waterfront Commission*.

([38]) Solove’s main concern with digital dossiers is that they could be abused by an incompetent or oppressive state (*The Digital Person*).

([39]) President Obama proposed a similar plan while this chapter was in press.