13. Bulk data collection, national security and ethics

Scott Robbins

1. INTRODUCTION

The rise of Internet communications has necessitated a rise in digital national security intelligence collection (including counter-terrorism intelligence and military intelligence) – currently at a scale never seen before in liberal democracies. The Snowden revelations of 2013 exposed digital intelligence collection that was pervasive and perhaps illegal (Greenwald 2013). People around the world were shocked at the capabilities of the National Security Agency (NSA), and the intelligence-collection practices revealed by Snowden have not slowed. On the contrary, many of these practices are being enshrined in the law (Pieters 2016; Travis 2016). Whether or not these practices are legal, it is essential to understand whether or not they are ethical – or how these practices can be conducted ethically. This involves identifying what makes these practices different from those that came before. Then, one must highlight how this changes the ethical analysis.

Two broad ethical paradigms constrain the practice of intelligence. First, there is what is acceptable for law enforcement – which generally takes a case-by-case approach to evaluating the acceptability of collecting intelligence or surveilling a subject or subjects. Essential considerations for law enforcement are: that there is reasonable suspicion or probable cause that the suspect (or suspects) are going to commit a serious crime, that the intrusion of their privacy is not disproportionate to the violations of citizens who will be the victims of that crime, and that the intelligence collection is necessary (that is, there is no less-intrusive alternative). Second, there is what is acceptable for national intelligence agencies to do during war, a context in which there are few constraints on their intelligence collection and analysis activities. Bulk data collection (BDC) for counter-terrorism purposes poses problems for each of these paradigms for two reasons. First, since terrorism is a crime, it can and should be dealt with as a crime by law enforcement (Miller 2008); however, terrorist groups, such as the Islamic State (ISIS), have at times launched military-style campaigns that target the state as a whole and, therefore, may require wartime tactics in response. Second, by definition, BDC sweeps up large amounts of data on innocent people, which is not something typically allowed by law enforcement. This has created a murky situation concerning counter-terrorism intelligence collection and analysis. This chapter cannot solve this complex problem or, rather, set of problems; however, it does provide some clarity on, and justification for, constraints that ought to be imposed on one specific form of intelligence collection: BDC.

Contemporary scholars have frequently discussed the ethics of intelligence activities within a Just War Theory (JWT) framework – those principles deemed necessary for the ethical initiation, conduct, and termination of war. Scholars have made efforts to modify JWT into a Just Intelligence Theory (JIT) (Bellaby 2012, 2016; Gendron 2005; Macnish 2014; Omand and Phythian 2013; Quinlan 2007). My focus in this chapter is to apply some of the latest work in JIT to BDC. Thus far, there has been no comprehensive ethical review of the practice of BDC for intelligence purposes.[1]

2. BULK DATA COLLECTION

To collect in bulk roughly means that the scope of collection will likely pick up many records that are not associated with current targets (Anderson 2016; National Research Council 2015). For example, the intelligence community (IC) may want all records associated with the current ISIS leader Abu Ibrahim al-Hashimi al-Qurashi. If the IC were only to collect records associated with him, then the IC is not collecting in bulk; instead, they are conducting a targeted collection. However, if the IC wants, for instance, all records coming into and out of Syria because they think many terrorists are operating there, then the IC is collecting in bulk. There are many Syrians whose data will be collected who are not engaged in terrorist acts and who do not even interact with terrorists. This is significant from an ethical standpoint because the IC is knowingly collecting data on innocent people and doing so on a large scale. This will be important for evaluating BDC in terms of the ethical principle of proportionality (see below).

BDC is done in two ways:

1. Bulk Interception: the practice of intercepting Internet communications data that are in transit.
2. Bulk Acquisition: the practice of acquiring bulk data from telecommunications and Internet companies.[2]

Bulk interception is accomplished by placing fibre optic splitters on telecommunications entry points. These fibre optic splitters copy the data and pass it along to intelligence agency infrastructure. These data will be filtered – to ensure that the data collection meets legal requirements[3] and that as little irrelevant data end up on agency servers as possible.

Bulk acquisition works in two ways. First, intelligence agencies can simply ask (or force) third-party institutions to turn over data in bulk (that is, data resulting from the application of some filter). Second, intelligence agencies may have back-door access to third-party institution servers. The Snowden revelations revealed that such back-door access was given to the NSA by Google, Facebook, and others (Greenwald and MacAskill 2013). In this chapter, BDC is also taken to be prima facie wrong, given that it involves infringing the privacy rights of innocent citizens on a large scale. The purpose of this chapter is to understand what conditions would have to be met for its use to be justified.

Privacy or other rights of any given *targeted* individual – that is, person who is an object of prior reasonable suspicion – cannot be the sole focus in the ethical evaluation of BDC. By definition, BDC is not targeted in this sense. Instead, it is the members of an entire group of people whose data will be collected to isolate members of the group for scrutiny. These groups are the result of filters being applied to the data passing through the Internet. The filters themselves, then, are where the focus should lie for an ethical evaluation of BDC. It is these filters that delimit the set of potential 'targets'. Few would have objected to a filter that selects all data related to Osama bin Laden. In the case of bulk collection, the filters are, by definition, much broader. These filters are what should be evaluated – in other words, this chapter focuses on understanding what might make the use of a particular filter morally justified and what might not.

3. JUST INTELLIGENCE

As already mentioned, a prominent theoretical perspective in the field of intelligence ethics advocates adapting JWT to evaluate intelligence collection and analysis. The primary reason for basing an ethics of intelligence on the ethics of war is that the conduct of both war and intelligence collection involves actions that are prima facie unethical. In war, you are killing people, destroying bridges and cities, holding people captive, and so on. All of these things are ethically bad. However, there are cases when such actions are necessary, proportionate, and morally justified. A country being invaded by another country should be able to defend itself – including shooting at their invaders. JWT outlines principles that are held to be necessary and sufficient to justify going to war (*jus ad bellum*) and to justify the conduct of that war once it is waged (*jus in bello*). Michael Quinlan (2007) argues that the practice of intelligence must also be justified and limited. In other words, there should be conditions that justify starting an intelligence program and limitations on how to conduct that intelligence program justly. Quinlan (2007) names these *jus ad intelligentiam* and *jus in intelligentia*.

The reason for using a theory based on JWT for intelligence collection is that intelligence collection involves harm and/or rights infringements that need further justification. Intelligence collection can involve listening in on private conversations, torture, deception, interception of communications, and so on. All of these actions would also be ethically disallowed under normal circumstances.

Harms from BDC can be divided into two types: privacy infringements and restrictions on autonomy. The data swept up by an intelligence agency belong to someone. An individual owns the information that those data reveal (Bellaby 2012). Prima facie, no one should be allowed to take these data. Of course, if this person is a known terrorist, then a state would be justified in collecting all information about this person. The point is that a state needs to justify its actions concerning BDC because harm or rights infringements are associated with such intelligence programs. If the state fails to justify such infringements, then violations have occurred.

People's autonomy – including citizens of the bulk-data-collecting state – can be restricted – intentionally or unintentionally – by BDC programs. Public knowledge of government BDC could affect innocent people's autonomy whether or not their data are collected. The so-called 'chilling effect' is when governmental regulation and policy not directed at certain activities deters individuals from carrying out protected activities (Robbins and Henschke 2017).

4. JUST BULK DATA COLLECTION

It is not the purpose of this chapter to justify the use of JIT; rather, it is to use principles of JIT to tease out ethical issues that arise due to the practice of BDC. In what follows, I use the JIT principles of just cause, proportionality, right intention, and proper authority to uncover issues that must be overcome to justify BDC's use.

4.1 Just Cause

What would be a just cause for intelligence collection? As counter-terrorism is the most salient reason given in recent times for BDC, this analysis will be restricted to cases involving terrorism.[4] At first glance, it is clear that counter-terrorism is a just cause for an intelligence operation. If terrorists are attempting to conduct attacks on citizens of a country, that country has just cause to collect intelligence that would prevent those attacks. Arguably, things might not be so simple for the reason that 'the general threat of terrorism, the so-called War on Terror, for example, is too indistinct to offer any specific just cause for an operation' (Bellaby 2016, p. 313).

Someone might claim that the IC has just cause to collect intelligence on everyone in the world to prevent unknown future threats from being realized. Since the IC does not know where the threats could come from in the future, no restriction on data collection would occur. This argument is spurious even if one is working with a reasonably broad definition of national security.

However, this is not a complete picture for two reasons. First, BDC occurs on a spectrum. At one end of the spectrum are practices of BDC that are unquestionably targeted; at the other end are practices of BDC that cannot be in any way considered targeted. At the end of the spectrum where the most-targeted practices are conducted, there might be practices such as collecting all the available data about the citizens who live in a small town known to be the home of some terrorists. At the other end of the spectrum might be practices such as collecting all the available data about United States citizens and anyone else who has entered the United States or who has communicated with anyone who lives in the United States. The justification for a particular instance of BDC will depend in part upon where it falls on this spectrum. Second, there is a conceptual issue regarding the point at which intelligence has been collected. On one account (further explained below), it seems as if the NSA, for example, collects most of the data travelling through the Internet as most of the world's data is routed at some point through the United States – although there are attempts to change this.[5] On the NSA's account of collection, the NSA collects a tiny fraction of the data travelling through the Internet. The result of this analysis will affect when the just cause principle can be applied.

Now comes the conceptual issue of what counts as 'collection', as it is not simple in the case of BDC. When can data be said to have been 'collected' by an intelligence agency? It may be helpful to take a rudimentary look at an email that ends up in the hands of an intelligence analyst through BDC:

When the email is sent, it gets routed to the backbone of the Internet run by (mostly) US communications companies (for example, AT&T).[6] The communications company acts as the post office in that it makes sure the communication is directed towards the intended recipient. It is here, at this first stage of the process (Stage 1), that, for example, the NSA has a splitter on the fibre optic cables to copy the data. At this stage of the process, the data would have to be stored until filters could be run on it. At the next stage of the process (Stage 2), the filters go through the data, discarding information that does not match any of the filters. At Stage 3, the data that make it through the filters end up on NSA servers for storage. Finally, at Stage 4, an analyst queries the data, resulting in the email (along with other data, perhaps) being returned to the analyst who reads it.

With Stage 1, above, it is clear that, for some time, the email is stored on a government server. NSA-owned equipment has possession of the data; however, at least as I have described the process, there is no potential for analysts to access those data.[7] An example from the physical world may help clarify the point. When you put your bag on the conveyor belt, it now sits on airport security property. If the machine that selects baggage for inspection were automated (with no human in control), this would be much like the BDC situation. All bags must pass through, but only a few are passed on for further inspection. We would hardly say that our bags have been collected (or that our privacy has been infringed) simply because they are on the conveyor belt. But once that bag is directed away from all the other bags towards the inspection team, the bag has been 'collected' (Stage 3). In this analogy, the bag going through the machine is like the data in temporary storage – it rests on the collector's property. Still, it is inaccessible to them (again, provided that the baggage machine is automated).

The intervention at Stage 2 appears trivial at first glance. It is merely the state of the data as filters are being run on them. It should look like a series of questions: Did this data come from Syria? No. Iraq? No. Is it encrypted using tools known to be used by terrorists? No. And so on. If any of the questions results in a yes, then the data move on to long-term or permanent storage. I include Stage 2 in my discussion because I want to highlight the difference between using these filters and running complex pattern-matching algorithms. Filters appear to be merely automating a human process. If one were to print out all of the emails passing through the Internet, a human could, in principle, check to see which of the emails matched one of the filters. Computers speed this process, but a human being could easily double-check each communication if need be. This is opposed to complex computer algorithms attempting to find patterns in the data or make predictions on the data. For example, a deep-learning algorithm could be trained on all of the communications associated with terrorism (previously) and used to classify future communications in terms of the connection with terrorist communications or other terrorist actions. This is no longer the automation of a human process; rather, it is a novel process that would be opaque to human minds. What can be said about algorithms like these being run on the data in temporary storage? Earlier I argued for evaluating the filter for just cause, but in this case, the filter is opaque to evaluation. The computer scientist who created the original algorithm would not even be able to explain how the algorithm classified a particular communication as being associated with terrorism. We would lack meaningful human control over how the algorithm selects communications to collect (Robbins 2019a; Santoni de Sio and van den Hoven 2018).

While the filter cannot be evaluated in the case of a machine-learning algorithm, an argument could be made that, if the algorithm is better at classifying communications in terms of their connection to terrorism than the articulable filters are, then the fact that they are not articulable should not be a reason not to use them. In other words, using machine-learning algorithms could be better for privacy because they are more accurate in their classifications. A similar point has been made in other contexts (Esteva et al. 2017; Robbins 2019b). This argument, however, fails in the context of counter-terrorism. First, the reason that the IC is collecting data in bulk is in part because of the changing communication tactics of terrorist groups. The classification of communications into those relevant to terrorist activity, and those not relevant, will change drastically over time. This is so for three reasons: first, as technology changes, the way we, as a society, communicate changes; second, terrorist groups of the future may communicate drastically differently than terrorist groups of the past; and third, terrorist groups know they are being surveilled and modify the way they communicate to thwart intelligence agencies. Therefore, it will be challenging to say that an algorithm is better at classifying communications than is an articulable filter.

At Stage 3, it is common to classify the data as collected. In this case, the data rest on government servers with access given to analysts under institutional constraints. These data are justifiably collected when there is evidence that there is a terrorist threat being organized or planned by the group described in the filter resulting in the collected data *and* that this threat is directed at the state collecting those data. While this may satisfy just cause, whether or not it is proportionate to the threat is another question.

4.2 Proportionality

Proportionality is a comparative notion where we judge that ‘an act is wrong if the relevant harm it will cause is out of proportion to its relevant good’ (Henschke 2018). Talk of proportionality with respect to going to war (*jus ad bellum*) is stated as a condition that ‘the destructiveness of war must not be out of proportion to the relevant good the war will do’ (Hurka 2005, p. 35). The principle of proportionality is also used for evaluating the just conduct of war (*jus in bello*), albeit in the context of the principle of discrimination and the principle of military necessity. According to the principle of necessity, the action must serve a military purpose. According to the *jus in bello* proportionality principle, the (unintended) deaths of innocent civilians, while permissible if militarily necessary, must not be disproportionate in the sense that the number of innocent deaths is disproportionate relative to the importance of the military objective (Hurka 2005).

With BDC, one can quickly see that the evaluation of proportionality hinges on empirical data. For any proportionality calculation, ‘we need specific facts about the costs [and] we need specification about the ends [that] are being sought’ (Henschke 2018). The extent of the harm done by BDC is challenging to determine before it has been carried out. How pervasive is the chilling effect mentioned in Section 3 above? A Pew Research Center poll concluded that 25 per cent of Americans had changed their online behaviour due to perceived government surveillance (Gao 2015). Depending on the methods used, the harms could be even more widespread – and more difficult to quantify.

The bulk acquisition of data from third-party institutions – especially when it pertains to back-door access and data retention – could result in diminished trust in participating institutions. Edward Snowden, in an interview with *The New Yorker*, explicitly told people not to use Dropbox, Google, or Facebook because of their susceptibility to intelligence collection (*The New Yorker* 2014). This, in turn, could harm the profits of third-party institutions and the US economy itself.

It will be necessary going forward to understand the harms to third-party institutions as a result of BDC. Harms like these must be taken into account in any calculation of proportionality. These harms would then have to be weighed against the efficacy of the program – or the good that it will do, which of course, is another empirical matter.

4.3 Right Intention

If the government intends to prevent terrorism, then right intention should be of little concern. Much like just cause above, the situation is not so simple. There may be a clear threat in Afghanistan of terrorism directed at the United States – a threat that constitutes just cause for BDC. However, the intention of the collecting state may be to glean information helpful to influence elections there. If that were the case, then the collecting state does not meet right intention.

What complicates right intention, however, is when and how often it should be applied. Right intention should be applied to the decision to create a filter that results in BDC. However, there is a time dimension that complicates this in two ways: (1) the filter will continue to collect data long after the decision was made to use that filter, and (2) the collected data will be stored long after that decision.

To illustrate the problem with (1) above, let us act as if BDC was a tactic to combat the Irish Republican Army (IRA), and the British intelligence agencies had just cause to collect all of the data coming into and out of Ireland. The IRA is no longer the threat it once was, so not only would British intelligence have to re-evaluate just cause, but they may have a problem with right intention as the British intelligence agencies may leave the filter because the data could be useful in the future. The time dimension of BDC means that the collected data should be tied to the justification for the creation of the filter – and deleted when that justification no longer holds.

4.4 Proper Authority

One could go along with traditional JWT and claim that BDC's only proper authority is the state. If this is the case, then a problem arises because, in practice, there are many third-party institutions collecting data in bulk. The practice of bulk acquisition is about the state copying data that have already been collected by third-party institutions – either by request or by back-door access. The question becomes whether or not the third party is then collecting bulk data as part of an intelligence collection and analysis program.

In many instances, this is not the case at all. Telecommunications and Internet companies store a lot of data that are necessary to conduct their business. Google does not store your email on their servers for reasons of national security; they store your email so that you have access to it. There is nothing inherently wrong with the IC obtaining data from third parties. If Osama bin Laden had a Gmail account, it would, and should, be expected that the NSA ask Google for those records – and it would, and should, be expected that Google provides them.

Things get more interesting if we look at forced data retention policies – in which laws mandate that third-party institutions retain data they may not typically retain for counter-terrorism (or national security). Now, the third-party institution is engaging in BDC as an intelligence program. This fails the principle of proper authority. Not only this, but now all of the data that have been retained that usually would not be should be included in our evaluations of just cause, right intention, and proportionality.

This problem is exacerbated when it is understood what the broad purpose of retaining such data would be. The purpose is, purportedly, national security. So the government faces a dilemma concerning the value of these data. Either the data are essential for national security, or they are not. If the data are essential, then the storage of those data should not be contracted out to third-party institutions. This is both because of the security risk of third parties being hacked and the blurring of institutional aims that such storage causes. Blurring these institutional responsibilities could damage the company's reputation and make it easier for those wishing to evade detection to choose other institutions. If the data are not essential, they should not force third-party institutions to retain such data.

5. CONCLUSION

This chapter has used JIT to evaluate the practice of BDC in liberal democracies for intelligence purposes. JIT forced the selection of an object of evaluation for BDC – in this case, the filters used to funnel data into government servers – and teased out some important ethical issues surrounding the practice. Most importantly, this evaluation pointed us to some essential constraints that should be placed on this practice. These constraints included: not using artificial intelligence as filters; the group specified by a particular filter must pose a threat to the collecting state; collected data must be tied to a filter and deleted when the justification for that filter no longer holds; and consumer companies, such as Google and Facebook, should not be allowed to act as intelligence agencies (collect data for the sole purpose of counter-terrorism).

This evaluation is just a start; however, it points to constraints that are not currently in place on BDC. Furthermore, this chapter starts from the premise that BDC is a valuable tool in the fight against terrorism. This may not be the case. If this tool turns out to be ineffective, it should not be used with or without the constraints outlined above. The point is that if intelligence agencies want this tool in their arsenal, they should be using it in a way that conforms to liberal-democratic principles and values. Having a just cause and right intention to collect data in bulk that are proportional to the threat and conducted by a proper authority would be a good start.

NOTES

[1] Bellaby (2016) does give an in-depth ethical evaluation of cyber intelligence (broadly construed) with a couple of paragraphs on what he calls '*en masse* collection', the term he gives to BDC.

[2] Bulk interception and bulk acquisition are terms used by David Anderson (2016) in his review of the UK's proposed Bulk Powers Act, which later became the Investigatory Powers Act.

[3] In the United States, for example, there must be minimization procedures to ensure that as little US personal data as possible ends up on intelligence agency servers. See, for example, Blum (2008).

[4] However, this analysis will apply to any context where national security is at stake.

[5] See, for example, Edmundson et al. (2016).

[6] This is not always the case, and increasingly there are methods for preventing your messages from being routed through surveillance states. See Edmundson et al. (2016).

[7] Although if XKeyScore exists as described by *The Intercept* (Lee et al. 2015), then analysts *do* have access and the BDC program would fail to meet just cause.

REFERENCES

Anderson, David (2016), *Report of the Bulk Powers Review*, United Kingdom: Williams Lea Group, available at https://terrorismlegislationreviewer.independent.gov.uk/wp-content/uploads/2016/08/Bulk-Powers-Review-final-report.pdf.

Bellaby, Ross W. (2012), 'What's the Harm? The Ethics of Intelligence Collection', *Intelligence and National Security*, 27 (1), 93–111, https://doi.org/10.1080/02684527.2012.621600.

Bellaby, Ross W. (2016), 'Justifying Cyber-Intelligence?', *Journal of Military Ethics*, 15 (4), 299–319, https://doi.org/10.1080/15027570.2017.1284463.

Blum, Stephanie Cooper (2008), 'What Really Is at Stake with the FISA Amendments Act of 2008 and Ideas for Future Surveillance Reform', *Boston University Public Interest Law Journal*, 18, 269.

Edmundson, Anne, Roya Ensafi, Nick Feamster, and Jennifer Rexford (2016), 'Characterizing and Avoiding Routing Detours through Surveillance States', ArXiv: 1605.07685 [Cs], May, http://arxiv.org/abs/1605.07685.

Esteva, Andre, Brett Kuprel, Roberto A. Novoa, Justin Ko, Susan M. Swetter, Helen M. Blau, and Sebastian Thrun (2017), 'Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks', *Nature* 542 (7639), 115–18, https://doi.org/10.1038/nature21056.

Gao, George (2015), 'What Americans Think about NSA Surveillance, National Security and Privacy', Pew Research Center blog, 29 May, accessed 23 September 2019, at http://www.pewresearch.org/fact-tank/2015/05/29/what-americans-think-about-nsa-surveillance-national-security-and-privacy/.

Gendron, Angela (2005), 'Just War, Just Intelligence: An Ethical Framework for Foreign Espionage', *International Journal of Intelligence and CounterIntelligence*, 18 (3), 398–434, https://doi.org/10.1080/08850600590945399.

Greenwald, Glenn (2013), 'NSA Collecting Phone Records of Millions of Verizon Customers Daily', *The Guardian*, 6 June, accessed 23 May 2018, at https://www.theguardian.com/world/2013/jun/06/nsa-phone-records-verizon-court-order.

Greenwald, Glenn, and Ewen MacAskill (2013), 'NSA Prism Program Taps in to User Data of Apple, Google and Others', *The Guardian*, 7 June, accessed 23 May 2018, at https://www.theguardian.com/world/2013/jun/06/us-tech-giants-nsa-data.

Henschke, Adam (2018), *Ethics in an Age of Surveillance: Personal Information and Virtual Identities*, Cambridge, UK: Cambridge University Press.

Hurka, Thomas (2005), 'Proportionality in the Morality of War', *Philosophy & Public Affairs* 33 (1), 34–66, https://doi.org/10.1111/j.1088-4963.2005.00024.x.

Lee, Micah, Glenn Greenwald, and Morgan Marquis-Boire (2015), 'A Look at the Inner Workings of NSA's XKEYSCORE', *The Intercept*, 2 July, accessed 5 April 2021, at https://theintercept.com/2015/07/02/look-under-hood-xkeyscore/.

Macnish, Kevin (2014), 'Just Surveillance? Towards a Normative Theory of Surveillance', *Surveillance & Society* 12 (1), 142–53.

Miller, Seumas (2008), *Terrorism and Counter-Terrorism: Ethics and Liberal Democracy*, Hoboken, NJ: Wiley.

National Research Council (2015), *Bulk Collection of Signals Intelligence: Technical Options*, Washington, DC: The National Academies Press, https://doi.org/10.17226/19414.

Omand, David, and Mark Phythian (2013), ‘Ethics and Intelligence: A Debate’, *International Journal of Intelligence and CounterIntelligence* 26 (1), 38–63. https://doi.org/10.1080/08850607.2012.705186.

Pieters, Janene (2016), ‘Proposed Law Allows Massive Data Mining by Intelligence Agencies’, *NL Times*, 15 April, accessed 15 May 2018, at https://nltimes.nl/2016/04/15/proposed-law-allows-massive-data-mining-intelligence-agencies.

Quinlan, Michael (2007), ‘Just Intelligence: Prolegomena to an Ethical Theory’, *Intelligence and National Security* 22 (1), 1–13, https://doi.org/10.1080/02684520701200715.

Robbins, Scott (2019a), ‘AI and the Path to Envelopment: Knowledge as a First Step towards the Responsible Regulation and Use of AI-Powered Machines’, *AI & SOCIETY*, 35 (2), 391–400, https://doi.org/10.1007/s00146-019-00891-1.

Robbins, Scott (2019b), ‘A Misdirected Principle with a Catch: Explicability for AI’, *Minds and Machines*, 29 (4), 495–514, https://doi.org/10.1007/s11023-019-09509-3.

Robbins, Scott, and Adam Henschke (2017), ‘The Value of Transparency: Bulk Data and Authoritarianism’, *Surveillance & Society* 15 (3/4), 582–89, https://doi.org/10.24908/ss.v15i3/4.6606.

Santoni de Sio, Filippo, and Jeroen van den Hoven (2018), ‘Meaningful Human Control over Autonomous Systems: A Philosophical Account’, *Frontiers in Robotics and AI* 5, Article 15, https://doi.org/10.3389/frobt.2018.00015.

The New Yorker (2014), ‘The Virtual Interview: Edward Snowden – *The New Yorker* Festival’, accessed 23 September 2019, at https://www.youtube.com/watch?v=fidq3jow8bc.

Travis, Alan (2016), ‘“Snooper’s Charter” Bill Becomes Law, Extending UK State Surveillance’, *The Guardian*, 29 November, accessed 3 May 2018, at http://www.theguardian.com/world/2016/nov/29/snoopers-charter-bill-becomes-law-extending-uk-state-surveillance.