Leonhard Menges*

Did the NSA and GCHQ Diminish Our Privacy? What the Control Account Should Say

<https://doi.org/10.1515/mopp-2019-0063>

Abstract: A standard account of privacy says that it is essentially a kind of control over personal information. Many privacy scholars have argued against this claim by relying on so-called threatened loss cases. In these cases, personal information about an agent is easily available to another person, but not accessed. Critics contend that control accounts have the implausible implication that the privacy of the relevant agent is diminished in threatened loss cases. Recently, threatened loss cases have become important because Edward Snowden's revelation of how the NSA and GCHQ collected Internet and mobile phone data presents us with a gigantic, real-life threatened loss case. In this paper, I will defend the control account of privacy against the argument that is based on threatened loss cases. I will do so by developing a new version of the control account that implies that the agents' privacy is *not* diminished in threatened loss cases.

Keywords: privacy, control, Edward Snowden, state surveillance, leeway control, source control, Frankfurt cases, threatened loss cases

1 Introduction

Many privacy scholars argue that privacy is essentially a certain kind of control. What we worry about when we worry about privacy on the one hand and Google, Facebook, the NSA, and so on on the other, the idea goes, is that we may lose control over personal information, or over access to personal information. This idea is so popular that it is often presented as the standard account of privacy (see, e. g. Rickless 2007, p. 779; Barocas and Nissenbaum 2014, p. 45; Macnish 2018, p. 419).

However, there are debates about whether the control view really offers the best account of the concept of privacy. In this paper, I will focus on a particular family of cases that have been around for more than 40 years and are often

*Corresponding author: Leonhard Menges, Department of Philosophy (KGW), University of Salzburg, Salzburg, Austria, E-mail: leonhard.menges@sbg.ac.at
<https://orcid.org/0000-0002-0539-0787>

presented as challenges for control accounts. Following Parent (1983b, p. 344), I will call them *threatened loss cases*. In threatened loss cases, very roughly, personal information about an agent is readily and easily available to another person, but not accessed. Most authors contend that control accounts imply that the privacy of the relevant agent is diminished in threatened loss cases. And this is typically presented as a vice of the control account of privacy. In this paper I will challenge the first assumption. The main point I will argue for is that there is a plausible version of the control account of privacy that implies that the agents' privacy is *not* diminished in threatened loss cases.

Some may ask why this should be relevant. Discussing conceptual questions about privacy has fallen out of fashion (see, e.g. Solove 2008, chs. 2 & 3; Nissenbaum 2009, introduction; Allen 2013). Moreover, the way I will approach these questions will remind some of the seemingly endless debates in analytic philosophy about science-fiction thought experiments or farfetched situations like Gettier, fission, or Frankfurt cases. And this will, I predict, convince some that the project of this paper is pointless. However, Macnish (2018) has recently effectively made a case for the claim that there still is a need for discussing the concept or definition of privacy. And threatened loss cases play a prominent role in his line of thinking. Let me elaborate.

Edward Snowden's revelation of how the NSA and GCHQ collected Internet and mobile phone data presents us with a gigantic, real-life threatened loss case. A great deal of personal information was readily available to them, but, in most cases, they did not access the relevant information. Macnish (2018, p. 418) observes that there were two different public responses to Snowden's revelation with regard to privacy. Many contended that the practices of the secret services invaded and diminished people's privacy, while others, including then-NSA director Keith Alexander, said that no privacy was diminished. Macnish suggests that the best explanation of this disagreement is that the two parties understand the concept of privacy in different ways. Those who say that no privacy was diminished seem to assume that a person's privacy is only diminished if another person learns something about her or actually accesses information about her. They can say that the NSA and GCHQ do not invade your privacy as long as no one actually reads your emails, finds out with whom you communicated, where you were, and so on. Those who disagree seem to assume that your privacy can be diminished by another person if the latter diminishes the control you have over your information or over access to your information. They can say that the NSA and GCHQ diminished your privacy by making or having the relevant data and information readily available.

It is important to distinguish between two different questions with regard to the NSA and GCHQ cases. First, one can discuss the *normative* question of

whether their activities were legitimate. What we need in order to answer this question is an account of the values or norms disrespected or violated by the NSA and GCHQ and what has been gained by their activities. This is the most politically relevant project because it may help us to reform the laws that govern the secret services.

Second, one can discuss the *conceptual* question of whether privacy is such that the activities of these institutions diminished the privacy of the citizens involved. What we need to answer this question is an account of the concept of privacy. Such an account is only indirectly politically relevant. For example, it does not suggest how to reform laws. However, it may help to frame the normative and political debate in an adequate, clear, and fruitful way. When Keith Alexander and his opponents disagree about whether or not the NSA diminished citizens' privacy, finding out who is right requires a good understanding of the concept of privacy. Moreover, even if a conceptual account fails to frame actual political debates (which is, perhaps unfortunately, not improbable), a convincing account of a contested concept like privacy is itself philosophically interesting.

This paper primarily contributes to answering the conceptual question. I will discuss some normative issues in Sections 5 and 6, but the main goal is to understand the concept of privacy by defending a prominent account against a long-standing objection. The tradition of discussing conceptual questions by using thought experiments and counterexamples in analytic philosophy provides us with tools that may help us to better understand threatened loss cases and their role in exploring the concept of privacy. This is why I will look for inspiration to one of these seemingly endless debates about certain cases.

I will begin by presenting threatened loss cases in more detail (Section 2). Then, I will show that there are important similarities between threatened loss and so-called Frankfurt cases in the discussion about responsibility and determinism. I will argue that these similarities point to a new way to understand the control account of privacy (Sections 3 and 4). At this point, I will have shown that there is room for a control account which implies that people's privacy is not diminished in threatened loss cases. The next step is to argue that there is good reason to accept this version of the control account. I will present some of these reasons in Section 5. In Section 6, I will apply the resulting control account to the NSA and GCHQ case and draw some general conclusions.

Let me make two clarifying points before I present threatened loss cases. First, I will follow prominent opponents of the control account and primarily talk about privacy as a state in which people can be (Parent 1983b, 1983a; Allen 1988, chap. 1, 2013; Davis 2009; Nissenbaum 2009, p. 69ff.; Macnish 2018). And I assume that the right to privacy is, then, the right to be in the state identified by

the best account of the concept of privacy. This should be seen as an explicit *stipulation*. I am well aware that it touches on controversial issues, but this would also be true of any other stipulation about the relation between privacy and the right to privacy. And diving into this relation would lead us too far away from the main points of this paper.

Second, I will only focus on informational privacy, and I will not say anything about locational, bodily, decisional, or other kinds of privacy. This seems to be acceptable because those who take threatened loss cases to pose challenges for control accounts focus on the informational context (Inness 1992, chap. 4 takes threatened loss cases to support control accounts and she uses a broader notion of privacy). And the proposal that I will make here should be understood as a reply to this challenge.

2 Threatened loss cases

The first threatened loss case can be found in a footnote in Judith Jarvis Thomson's seminal paper on the right to privacy:

If my neighbor invents an X-ray device which enables him to look through walls, then I should imagine I thereby lose control over who can look at me: going home and closing the doors no longer suffices to prevent others from doing so. But my right to privacy is not violated until my neighbor actually does train the device on the wall of my house (Thomson 1975, n. 1).

Thomson discusses the right to privacy, not the state of having privacy, on which I focus in this paper. William Parent also focuses on the state of privacy and slightly modifies Thomson's original case:

Suppose *A* invents a fantastic X-ray device that enables him to look right through walls. *A* then focuses the device on my home but refuses to use it. Since he certainly has the power to find out everything that I am doing in my home it cannot be said that I any longer enjoy control over personal information about myself vis-à-vis *A* – at least I don't in regard to activities done at his home. Still *A* has not invaded my privacy. He doesn't do that until he actually looks through his device. So while the lack of control certainly threatens privacy it does not necessarily involve its loss (Parent 1983b, p. 344; see also Rickless 2007, p. 783; Davis 2009, p. 457; Moore 2010, p. 21 n. 35).

Parent and his followers present this case as a challenge for the control account of privacy. The basic argument is that control accounts imply that my privacy is diminished in this case, that this is intuitively implausible, and that, therefore,

control accounts are intuitively implausible. Kevin Macnish argues for the same conclusion by relying on a different threatened loss case:

[I]magine that I leave my diary on a table in a coffee shop and return to that shop 30 minutes later to retrieve it. When I enter the shop I see a stranger with my diary on her table, a different table from the one at which I was sitting. I therefore know that she, or someone, has moved my diary; but have they read it? I have not been in control of my diary for half an hour, in which time anything might have happened to it (Macnish 2018, p. 420).

Macnish (2018, p. 425) suggests that, in a version of this case, in which the stranger has not opened the diary, control but no privacy has been diminished, such that losing control is not sufficient for losing privacy. He also contends that this version is, with one important difference, relevantly similar to the real life situation in which many of us are with respect to the NSA and GCHQ. The important common point is that personal information is readily available to the NSA, GCHQ, and the finder of the diary, but in most cases the information has not been accessed. The difference is that the secret services ‘have actively pursued the collection of data’ (Macnish 2018, p. 426), while in the case sketched above, the other person finds the diary in a café.

Let us take a step back from the particular scenarios. Then we can see that threatened loss cases have a common structure. Personal information about an agent *A* is readily available to another agent *B*, *B* knows about this, can easily intervene and, thereby, access and learn the information about *A*, yet *B* does not do so. There seems to be consensus in the literature that the fact that *B* can easily access and learn personal information about *A* and that *B* knows that this is so diminishes *A*’s control over the information. Moreover, it is taken for granted that the control account of privacy implies that this also diminishes *A*’s privacy. For many, this is an implausible implication and they conclude that the control account should be dismissed.

Proponents of control accounts can reply in two ways. First, they can argue that it is, contrary to what has just been said, intuitively plausible that privacy is diminished in the cases presented above. That is, they could try to turn the tables and say that it is a virtue rather than a vice of control accounts that they imply that privacy is diminished in the diary, X-ray and NSA cases. I will briefly come back to this line of defense at the end of Section 5. However, I share the intuition of those who find it more plausible that privacy is not diminished in threatened loss cases. Therefore, I will opt for the second reply and argue that, contrary to what has been said above, the control account does not necessarily imply that privacy is diminished in threatened loss cases. This is what I will show in the next section.

3 From Frankfurt to threatened loss cases

In what follows I will argue for two claims. First, we are, in everyday life and philosophical discourse, familiar with a kind of control which is not diminished in threatened loss cases; second, proponents of the control account can use this notion in order to spell out their view that privacy is essentially about having control. I will proceed by developing an argument from analogy. I will show that threatened loss cases have striking similarities with Frankfurt cases (see Frankfurt 1969) in the debate about responsibility and determinism and I will argue that these similarities point to a new way to understand control accounts of privacy.

Let me begin with a bit of background about Frankfurt cases. The uncontested starting point is the idea that an agent is only responsible for a certain action if she has control over it. If you have no control over how your body moves, then you are not responsible for this movement. But what is this kind of control? Everyday thinking suggests that we should understand it as the ability to do otherwise, which I will call *leeway control* (see, e.g. McKenna and Pereboom 2016, ch. 2). If you step on my foot but you could not have done otherwise, then it seems natural to think that you are not responsible for stepping on my foot. Moreover, it is suggesting to think that determinism rules out a person's ability to do otherwise. Combining these ideas yields incompatibilism about determinism and responsibility. If leeway control is necessary for responsibility and if determinism rules out leeway control, then nobody is responsible in a deterministic world.

Frankfurt cases aim at showing that the first assumption is false: a person can be responsible for what she does even if she cannot do otherwise. There are two main ideas associated with Frankfurt cases (see Sartorio 2016 for an overview). First, we should distinguish between factors that explain why an agent acts the way she does and factors that explain why she cannot act otherwise. Second, the latter factors by themselves do not undermine the responsibility of the agent. In most everyday cases both factors fall together. If I step on your foot because someone pushes me, then the factor that explains why I stepped on your foot is the very same factor that explains why I could not have done otherwise: namely because someone pushed me. But there are cases in which the two factors fall apart – Frankfurt cases. In these cases, there is a factor that explains why an agent has no leeway control (cannot do otherwise), but this factor does not explain why the agent acts the way she does. In these cases, it seems that the agent is responsible for what she does. Thus, Frankfurt cases aim at showing that one can have the control which is necessary for being responsible for an action even if one does not have leeway control over it.

Here is a typical Frankfurt case (from the good old days):

Because he dares to hope that the Democrats finally have a good chance of winning the White House, the benevolent but elderly neurosurgeon, Black, has come out of retirement to participate in yet another philosophical example. [...] He has secretly inserted a chip in Jones's brain that enables Black to monitor and control Jones's activities. Black can exercise this control through a sophisticated computer that he has programmed so that, among other things, it monitors Jones's voting behavior. If Jones were to show any inclination to vote for McCain (or, let us say, anyone other than Obama), then the computer, through the chip in Jones's brain, would intervene to assure that he actually decides to vote for Obama and does so vote. But if Jones decides on his own to vote for Obama (as Black, the old progressive would prefer), the computer does nothing but continue to monitor – without affecting – the goings-on in Jones's head.

Now suppose that Jones decides to vote for Obama on his own, just as he would have if Black had not inserted the chip in his head. It seems, upon first thinking about this case, that Jones can be held morally responsible for his choice and act of voting for Obama, although he could not have chosen otherwise and he could not have done otherwise (Fischer 2010, p. 316).

The factors that explain why Jones has no leeway control over his voting (such that he cannot refrain from voting for Obama) are Black and his fantastic devices. But the factors that explain Jones' actual voting for Obama are not Black and his devices, but some features of Jones himself, namely his beliefs, desires, intentions, and so on. Even if Black were not present, he would have voted for Obama and the process that led to his voting for Obama would have been exactly the same. Therefore, his actual voting for Obama still seems to be up to him in a certain sense and he still seems to be responsible for voting for Obama. I will call the kind of control that Jones has over his voting for Obama *source control* (see, again, McKenna and Pereboom 2016, chap. 2). The idea is, roughly, that we do not need the ability to do otherwise in order to be responsible for our actions. What we need is to be the right kind of source of our actions, such that the actions are *our* actions in a specific sense. I will come back to the question of how to spell out what it means to be the right kind of source below.

As many will know, the debate about Frankfurt cases goes on and on. And my aim is not to say anything illuminating about them here. My aim is rather to show that control theorists of privacy can find inspiration for a new account of threatened loss cases in the discussion about Frankfurt cases.

In what follows, I will focus on Macnish's (2018, p. 420) diary case and add a few lines in order to make the analogy clear.

Diary Case 2: Imagine that I leave my diary on a table in a coffee shop and return to that shop 30 minutes later to retrieve it. When I enter the shop I see a stranger with my diary on

her table, a different table from the one at which I was sitting. I therefore know that she, or someone, has moved my diary; but have they read it? Imagine that the stranger has not yet read it but wants to know what my last entry says. She has firmly decided to read it before 3 pm and she would read it even in my presence (imagine that she is very strong and I would not be able to prevent her from reading it). I come back at 2.55 pm and tell her: 'It's terrible, I'm forgetting everything these days! I hope I'm not getting ill. Actually, I wrote about it in my diary this morning. Please, look at the last pages.' In response to this, the stranger reads my last entry in the diary.

In order to see the structural similarities between Frankfurt and threatened loss cases, it is helpful to, first, think of Jones in the Frankfurt case as analogous to me in the new diary case – I will call them (Jones and me) agent *A*. Second, let us think of Black and his devices in the Frankfurt case as analogous to the stranger in the diary case 2 – agent *B*. Third, the event of Jones' voting for Obama corresponds to the event of the stranger's learning about my last entry in the diary – event *E*. Now, control theorists can argue in the following, to-be-specified way. In Frankfurt and threatened loss cases, *B* makes it the case that *A* does not have a certain kind of control over *E*. More specifically, *A* cannot effectively choose whether or not *E* happens. However, *A* still has another kind of control over *E*, namely *A* is the right kind of source of *E*'s actual taking place. And it is, control theorists of privacy can continue, the latter kind of control that is relevant for privacy. The conclusion is, then, that *A* still has privacy in this threatened loss case, despite losing leeway control, because *A* maintains source control. Let me elaborate.

In both cases, there are factors that explain why an agent *A* (Jones and I) cannot effectively choose whether or not event *E* happens (Jones votes for Obama, the stranger learns about my last entry). It is important to stress the 'whether or not'. For *A*, there is only one option to realize. Jones cannot effectively choose whether or not he will vote for Obama because if he shows any inclination not to, Black will make him vote for Obama. And I cannot effectively choose whether or not the stranger will learn about the last entry in my diary because if I were to ask her to not read it, she would read it anyway. That is, in both cases *A* does not have leeway control over whether or not event *E* happens.

The factors that explain why *E* happens have nothing to do with the factors that explain why *A* has no effective choice. The factors that explain why *A* has no effective choice are about another agent *B* (Black and the stranger). In the Frankfurt case, the factors that explain why Jones has no effective choice are Black and his devices. In diary case 2, the factors that explain why I cannot effectively choose whether or not the stranger learns about my last entry are that the diary is close to her on her table and that she has firmly decided to read it

before 3 pm (and that she is very strong). However, these factors about *B* do not explain why event *E* takes place. The factors that explain why *E* happens are about *A*, not about *B*. That is, Jones votes for Obama simply because he wants to and decides to do it. He would have done the same if Black and his devices were absent. I let the stranger read the entry because I want her to read it and decide to ask her to read it. I would have let her read it even if she had not decided to read it anyway (or if she wasn't that strong).

In an intuitive sense, event *E* is up to *A*. Jones' voting for Obama and the stranger's learning personal information about me are up to Jones and me; we make these things happen; we have an important kind of control (source control) over them.

Proponents of the control account of privacy can use this analogy by arguing that privacy should be understood as source control over (access to) information, not as leeway control. In threatened loss cases, agents lose leeway control over (access to) the relevant information. An agent *A* has this kind of control just in case she can effectively choose between different options, namely between whether *B* accesses or learns about the information or not. However, *A* still has source control over (access to) the information in threatened loss cases. Intuitively, we can think of source control as one's being able to knock on one's neighbor's door and tell her something new about oneself. And we can have this kind of control even if our neighbor would learn the information anyway. More formally, *A* has source control over (access to) information just in case *A* is the right kind of source of the information flow, if the information flows to another person at all.

When is an agent the 'right kind of source' of an information flow? Giving a full answer to this question would require, plausibly, a whole book. But there are two places where the control theorist can look for models. The first is compatibilism about responsibility and determinism (see, e. g. Frankfurt 1971; Fischer and Ravizza 1998). Proponents of this view look for an account of being the right kind of source of an action that supports the idea that an agent can be responsible for it even if she cannot do otherwise. One famous account says, very roughly, that Jones is the right kind of source if he does not only desire to vote for Obama but also desires to desire that he vote for Obama (see Frankfurt 1971). Control theorists about privacy could adopt this idea. They could say that I have privacy with regard to certain pieces of information just in case if these pieces of information flow to another person at all, then this is the result of my desiring it and my desiring that I desire it to flow in this way. In the diary case, the stranger would diminish my privacy if she learns about the last entry even though my first- or second-order desires were against this flow of information.

The second place to look for models for how to understand being the right kind of source of information flow is the debate about informed consent. The

idea is that an agent's consenting to a medical treatment, sexual intercourse, and so on, is only valid if the agent fulfills certain conditions. These conditions are often spelled out without requiring the straightforward ability to do otherwise. One prominent account says, for example, that the agent has to consent intentionally, she has to understand what is at issue, and she has to consent voluntarily (without being coerced, manipulated, and so on) (see, e.g. Beauchamp 2010; Beauchamp and Childress 2013, ch. 4). Adapting this view to a source control account of privacy would yield the following rough picture: I have privacy with regard to certain information just in case if the information flows to another person at all, then this is because I intend it to, I voluntarily let it flow in this way and I understand what is at issue. In a version of the case in which the stranger reads in the diary before I come back, these conditions are not fulfilled. I do not intend to let the information flow in this way and, thereby, my privacy is diminished.

I am well aware that these pictures of responsibility and of valid consent are highly contested. And I am not claiming that these are the *best* ways to make sense of what it is to be the right kind of source of an information flow. The important ideas that I want to make clear at this point are, first, that there is room for a kind of control over (access to) information which does not involve leeway control and, second, what the resulting control account of privacy *could* look like.

To sum up, just as responsibility theorists distinguish between leeway and source control over actions, privacy scholars should distinguish between leeway and source control over (access to) information. Frankfurt cases are the rare cases in which it becomes clear that agents do not have leeway but do have source control over what they do. Similarly, privacy scholars can say that threatened loss cases are such that it becomes clear that an agent can lose leeway control over (access to) personal information but maintain source control. Privacy scholars can add that the kind of control that is relevant for privacy is source control, not leeway control. According to the resulting source control account of privacy, the nature of privacy is being the right kind of source of information flows, if it flows at all. As this kind of control is not diminished in threatened loss cases, proponents of this account would say that privacy is not diminished in threatened loss cases.

4 Three views on privacy

One may ask now how exactly the source control account of privacy is supposed to differ from the leeway control and the access-based view. More specifically,

does it have implications that deviate in relevant ways from the other two accounts? In order to answer this question, I will apply the views to different cases.

Recall, first, the

Original Diary Case (see Macnish 2018, p. 420, p. 425): Imagine that I leave my diary on a table in a coffee shop and return to that shop 30 minutes later to retrieve it. When I enter the shop I see a stranger with my diary on her table, a different table from the one at which I was sitting. I therefore know that she, or someone, has moved my diary; but have they read it? Imagine that the stranger has not read it and has firmly decided not to read it. The stranger gives me the diary. I can tell the stranger about the content, but I don't.

This is a standard threatened loss case. Leeway control accounts imply that my privacy was diminished at the moment in which I lost the ability to effectively choose whether or not the stranger or someone else would read my diary. Source control and access-based views, by contrast, say that no privacy is diminished in this case because the stranger has not accessed the information and the relevant information has not flowed to another person.

Now consider

Diary Case 3: The only difference to the original case is that the stranger, unbeknownst to me, reads my diary before I come back. That is, I can tell the stranger about the content of my diary, though I don't, but the stranger has already learned about it anyway.

Again, leeway control theorists say that privacy has been diminished before the stranger reads my diary, namely when I lost effective choice with regard to the relevant information. By contrast, access-based and source control views say that I lose my privacy when the stranger actually reads my diary. However, the two accounts differ with regard to the explanation of why my privacy is diminished in this moment. Access-based views say that privacy is lost because the stranger actually accesses the relevant information. Source control views explain the loss of privacy by referring to the fact that the stranger lets information about me flow without my being the right kind of source of this flow.

Finally, take

Diary Case 4: This case is identical with case 2 from the previous section (the one analogous to the Frankfurt case) with the only difference being that the stranger does not intend to read my diary. That is, when I come back to the coffee shop, I freely and knowingly ask the stranger, who has not read my diary and does not plan to read it, to read the last entry of my diary. In response, the stranger reads it.

Again, leeway views say that I lost privacy when I lost effective choice. Access-based views say that I lose my privacy when the stranger reads the diary and,

thereby, accesses personal information. On the source control view, by contrast, no privacy is diminished. The view says that privacy is only diminished when information flows without one's being the right kind of source. Plausibly, however, I am the right kind of source of the information flow in diary case 4 because the stranger reads the diary in response to my giving valid consent.

To sum up, the three views discussed so far have different results when applied to specific cases. They differ with regard to how they answer the questions of whether privacy is diminished at all, of when, and of why privacy is diminished. Thus, the source control account is a substantial alternative to the leeway control and the access-based view.

5 Why opt for source control?

Recall that the main aim of this paper is to defend the control account against the objection that it has the wrong implication when applied to threatened loss cases. In order to achieve this aim, I have developed a new version of the control account and shown how it differs from standard alternatives. But are there reasons for accepting this view rather than one of the others? In what follows, I will mainly compare the source control with the leeway control account, but I will begin by giving a brief sketch of what I take to be an advantage of the source control over the access-based view (thanks to an anonymous reviewer for pressing me on this).

Recall that in diary case 4, the access-based view says that my privacy is diminished when the stranger reads my diary as a response to my giving valid consent to her reading the last entry. Similarly, the view implies that we lose privacy when we freely and knowingly tell our friends about our problems, desires, and secrets. According to this view, whenever someone else accesses certain information about us, our privacy is diminished.

I take this to be implausible (see also Inness 1992, p. 46). Intuitively, sharing private information with someone close to us does not necessarily involve *losing* or *diminishing* privacy. We do not have less privacy when we have shared privacy. But this is what the access-based view implies. The source control account, by contrast, says that when we are the right kind of source of the information flow, say, when we freely and knowingly tell our friends about something personal, then we do not lose privacy and no privacy is diminished. This seems to be the correct result.

Admittedly however, intuitions seem to clash here. Macnish, for example, explicitly says (2018, p. 422) that telling friends about personal matters diminishes

one's privacy. While he finds this intuitive, I find it very counter-intuitive. This suggests that more needs to be done to settle the dispute between access-based and source control views. In the meantime, I propose to draw the moderate conclusion that those who have the intuition that sharing private information does not necessarily diminish privacy should opt for the source control rather than the access-based view.

Let me now turn to a comparison of the leeway and the source control account. The first advantage of the source control account is obvious. Opting for this view avoids what many take to be an implausible implication of leeway accounts, namely that people lose privacy in threatened loss cases. As I have shown above, many authors reject control accounts at least partly because of this implication (e. g. Parent 1983b; Rickless 2007; Davis 2009; Macnish 2018). As the source control account does not have this implication, this is a good reason to adopt the source control account, rather than the leeway account.

The second advantage of the source control account is that leeway control accounts seem to have counter-intuitive implications if our world turns out to be deterministic. Leeway control accounts say that privacy involves the ability to effectively choose between different options with regard to who knows or accesses personal information. If determinism is true, then it seems as if nobody can ever make such an effective choice (see McKenna and Pereboom 2016, chap. 1 for an overview). Thus, it seems that, on this picture, no human has privacy in a deterministic world. This is a surprising implication. If your diary is in your safe and nobody but you has access to it and nobody knows about it as long as you don't consent to her learning about it, then the contents of your diary clearly are private. And this seems to be so, regardless of the truth or falsehood of determinism. The truth or falsehood of determinism seems to be irrelevant when we want to know whether or not a person has privacy with regard to certain information. However, according to leeway control accounts of privacy, the truth or falsehood of determinism seems to be relevant. And this speaks against the leeway control account.

The source control account has a better chance of avoiding the implication that nobody has privacy in a deterministic world. Both sketches of being the right kind of source of information flow that I have presented in the preceding section are compatible with the truth and with the falsehood of determinism. Even if determinism is true, you can desire and desire to desire to let another person know about your health condition and you can tell her about it voluntarily, intentionally and with a sufficient understanding of the situation. Thus, source control over (access to) information is, plausibly, compatible with determinism such that this view avoids the implication that no human has privacy in a deterministic world.

As an intermediate conclusion, the source control account avoids some counter-intuitive implications of the leeway control account of privacy. Proponents of the leeway control account may reply that this is only a half-hearted argument for the source control account. Leeway control over (access to) information is, they may say, what we actually care about, what is important to us, and what we really need. Thus, they could conclude that the leeway control account is the only view that makes sense of the intuitive value and importance of privacy.

As a reply, let me present some considerations that are typically taken to show that privacy as control in general is important and let me show that they also support the idea that privacy as source control is valuable. To avoid misunderstandings, the thesis I am going to argue for in the remainder of this section is *not* that the source control account is to be preferred over the leeway control account. More moderately, the goal right now is to show that the leeway control and the source control account can make sense of the intuitive value of privacy in very similar ways. However, this result may be combined with the claim defended above, which held that the source control account deals better with both threatened loss cases and the possible truth of determinism; and this yields a powerful argument to the conclusion that the source control account is to be preferred over the leeway control account.

First, being able to control who knows what about us is very helpful in a competitive and less-than-fully-tolerant society. It is important to control whether our employer knows what diseases we may have and in many situations it is important to control who knows our religion, sexual orientation, the family planning choices we may have made, and so on. In a world such as ours, controlling information about us protects us from an intolerant society and it protects our bargaining power.

Can the source control account make sense of this idea? Surely it can. What is valuable, according to this view, is being the right kind of source of information flow. For example, the source control account says that I have privacy with regard to what my boss knows about my health condition, as long as, if the relevant information flows to her at all, I am the right kind of source of this flow. And as long as this is so, my bargaining power is protected. But if someone tells her about my health condition without my being the source of this information flow, then I lose source control, and, thereby, my bargaining power is diminished. Thus, source control protects us and our bargaining power from an intolerant society.

Second, control over (access to) certain information about us protects our individual liberty and autonomy, and, thereby, our democratic system. There are many things that we would not do if we had no control over who will or can

learn about our doing it. These include trivial acts such as singing in the shower, but also discussing controversial political or religious questions. And being able to do these things is important for shaping our lives according to our own values (see, e. g. Fried 1968; Rössler 2004, chap. 3, 2017). This may be taken to suggest that privacy protects autonomy by detaching the individual from society. But, as many recent scholars have argued, the role of privacy need not be understood in this way (see, e. g. the essays in Rössler and Mokrosinska 2015). One can think of autonomy as being socially embedded such that it only develops and is exercised in relationships and in society more generally. And control over (access to) personal information is, as I will show in a moment, very helpful or even required for shaping relationships, and, thereby, for developing and shaping autonomy. Additionally, a well-functioning democratic system needs citizens who act autonomously in the public sphere. But it is very plausible that humans are such that they sometimes need a break from this sphere. We could not play our public roles if we had to play them all the time. Thus, we sometimes need control over who can observe us when we are not playing our public roles in order to, at other times, participate in public democratic processes (see, e. g. Rössler 2017; for similar ideas see Lever 2013; Roberts 2015; Mokrosinska 2018). Thus, control over (access to) information protects our individual liberty, our ability to live an autonomous life, and, thereby, our democratic system.

Again, source control over information can fulfill these important functions. The source control view says that we have privacy with regard to our controversial political views or with regard to what we do when we leave the public sphere as long as, if this information flows at all to other people, then we are the right kind of source of this information flow. If we fear that the information flows even though we do not play the right kind of role in this process, then we will not feel free to openly discuss our ideas, to experiment with ways of living, relationships, and so on. This will destroy the social context that we need to develop and exercise autonomy, it will restrict our individual liberty and it will not give us the break we need in order to participate in public life. Having source control would protect us from all this. Thus, source control over (access to) information protects autonomy, the social context in which it is exercised, individual liberty, and the democratic system.

Third, different relationships are partly constituted by what one reveals to the other party. Most friendships, for example, involve letting the friend know about certain aspects of oneself that one does not reveal to most others. If we have no control over information about us, it would be much harder for us to form relationships of different degrees of intimacy and closeness. And being able to distinguish between a close friend and a mere acquaintance is an important aspect of the well-being of most people. Thus control over (access

to) information about us protects an important aspect of the well-being of most people (see, e.g. Fried 1968; Rachels 1975; Gerstein 1978; Moore 2010, ch. 2; Marmor 2015).

Source control can serve this function, too. Our telling a friend about our health problem can be an expression of trust and intimacy, and, thereby, shape our friendship. If we are not the source of this information flow because someone else tells our friend about our health then this will not be an expression of our trust and intimacy and will not shape our relationship in the corresponding way. What we need in order to form relationships is source control over information flow. Thus, source control protects our ability to shape relationships.

To sum up, many control theorists have argued that control over (access to) personal information plays very valuable roles and that, therefore, the control account can make sense of the idea that privacy is important. I have argued that the source control account of privacy can adopt the same lines of thinking because source control over (access to) personal information plays the roles that control theorists want control to play.

Some control theorists will not accept this conclusion. So far, I have only considered cases in which source and leeway control views converge to the implication that privacy is diminished and that something valuable has been lost. However, source and leeway control accounts come apart, namely in threatened loss cases. As I said at the end of Section 2, some control theorists will say that it is plausible that privacy is diminished in these cases. Moreover, they may add that something valuable has been diminished. They could argue, for example, that privacy is important because it prevents others from obtaining power over us. In some threatened loss cases personal information about us is easily available to people who should not have power over us such that this protection has been lost (thanks to an anonymous reviewer for this point). Others may argue that we need leeway control over concealment and exposure in complex societies like ours. This valuable kind of control gets lost in threatened loss cases, too (thanks to another anonymous reviewer for this point). Based on these ideas, some may object that the source control account is essentially incomplete: because the source control view cannot explain that privacy and some values protected by privacy can get lost in threatened loss cases, we should dismiss it.

As a first reply, and as I will show in more detail in the next section, proponents of the source control view can make sense of the idea that important *values*, like protection from power and leeway control over concealment and exposure, get lost in threatened loss cases. They only insist that *privacy* does not get lost. Secondly, it seems that we have reached a clash of intuitions again. Whereas some proponents of leeway control views may find it intuitive that

privacy is diminished in threatened loss cases, many others (e. g. Parent 1983b; Rickless 2007; Davis 2009; Macnish 2018), me including, find this counter-intuitive.

In light of this apparent clash of intuitions, the main theses of this paper can be understood as being the following: first, it is possible to combine the intuitively plausible (but not universally accepted) assumption that privacy is not diminished in threatened loss cases with the idea that privacy is essentially a kind of control. Second, the resulting source control view has some interesting and plausible features.

6 Back to state surveillance

Recall that discussing the concept or definition of privacy has become important in recent years because of two opposed evaluations of state mass surveillance. While one group, including then-NSA director Keith Alexander, said that no privacy was diminished by what the NSA and GCHQ did, many others contended that privacy was in fact diminished. Let me apply the account I have developed to this debate.

Macnish interprets the NSA and GCHQ case analogously to the original diary case which is, then, analogous to Parent's classical X-ray case. And most authors in this debate agree that control accounts imply that privacy is diminished in these cases. Thus, they take the control account to be committed to the claim that the NSA's and GCHQ's making personal data readily available is a form of privacy invasion and diminishes citizens' privacy even if nobody actually accesses and learns the relevant information.

I have argued that this reading of control accounts presupposes a leeway control account of privacy, according to which privacy involves the ability to effectively choose whether or not others learn or have access to certain information. Once the relevant data are available to the NSA and GCHQ, you cannot do anything in order to stop them from accessing or learning the information. Thus, your leeway control is diminished.

I have also argued that there is an alternative to the leeway control account, namely the source control account of privacy. According to the source control account, privacy essentially consists in being the right kind of source of information flow to another agent if the information flows at all. In threatened loss cases, including the NSA and GCHQ case, the information does not flow to another agent as long as nobody actually accesses the data and learns something about the relevant citizens. To illustrate, imagine that the data collected by

the NSA would tell them that Anna wrote her psychotherapist an email on Friday at noon. As long as nobody has actually accessed this piece of information, Anna can still go to the NSA and tell everybody that she wrote her psychotherapist an email on Friday at noon. Thus, she can still be the right kind of source of this information flow. This kind of control is diminished, however, as soon as an NSA employee accesses the data before Anna tells her about it. Thus, the source control account of privacy says that Keith Alexander was right: no privacy is diminished as long as the information is not accessed.

In Section 1, I distinguished between the conceptual question of whether privacy is such that the activities of the NSA and GCHQ diminished citizens' privacy and the normative question of whether these activities are legitimate. So far, I have presented an answer to the conceptual question. One may now ask if this answer has implications for the normative project. One may wonder, for example, whether the claim that Keith Alexander was right in saying that the NSA did not diminish privacy as long as no information was accessed implies that there is nothing problematic about what the NSA did. Importantly, this does not follow. Let me elaborate.

Macnish (2018, Secs. 3 & 4) argues convincingly for the claim that what the NSA and GCHQ did may have been illegitimate even if no privacy was diminished. The first reason is that making someone believe that her privacy has been invaded can be bad for her even if her privacy was not in fact diminished. This has been well known ever since Bentham's theory of the panopticon in which the prisoners never know whether they are being observed and, therefore, always (or mostly) act as if they were being observed. This is a way of controlling citizens that diminishes their liberty to a problematic degree. Thus, letting citizens believe that their privacy can be diminished at any time is problematic in this respect.

The second line of thinking that shows that the activities of the NSA and GCHQ may be illegitimate is that what they did increases the risk that privacy will in fact be diminished. They have made the data and information readily available and, therefore, the risk that they will actually invade our privacy has become much greater. Making people subject to such a risk is problematic.

Third, an idea not discussed by Macnish and hinted at above is that it is important for citizens to be protected from the power of other agents, be it institutions or individuals. Making data about citizens available to these agents diminishes this protection. These agents can obtain power by accessing the data directly or by letting machines analyze them, for example in order to use the results for targeted political campaigning. Thus, the activities of the NSA and GCHQ may have been illegitimate because they diminished an important protection from power.

Macnish draws a general conclusion:

[F]ocusing on privacy distracts all involved from the real issues of harm to the general populace. Indeed, by focusing the argument on privacy it may be harder to persuade supporters of actions taken by NSA and GCHQ that what they are doing is harmful. They have a response to that argument, namely that they are not violating people's privacy except in specific, justifiable, targeted cases (Macnish 2018, p. 429).

I agree with this conclusion. What is normatively and politically problematic in the NSA and GCHQ case is not that people's privacy is diminished but other harms and wrongdoings. The important difference between Macnish's and my account is, however, that I take this conclusion to be compatible with and even supported by a control account of privacy, namely the source control account.

To sum up, I have argued for a new version of the control account of privacy, according to which having privacy with regard to information is being the right kind of source of information flow, if information flows at all. This view has many of the positive aspects of classical control accounts, offers a new interpretation of threatened loss cases and, thereby, avoids an implication of typical control accounts that many find problematic. Moreover, this view may help us to focus better on the severe harms and wrongdoings associated with state mass surveillance.

Acknowledgements: I am grateful to Peter Königs, Kevin Macnish, and two anonymous referees for helpful comments and suggestions. Thanks to Claire Davis for proof-reading.

References

- Allen, A.L. (1988). *Uneasy Access: Privacy for Women in a Free Society* (Totawa, New Jersey: Rowman & Littlefield).
- Allen, A.L. (2013). 'Our Privacy Rights and Responsibilities: Replies to Critics', *Newsletter of the American Philosophical Association: Philosophy and Law* 13 (1): 19–27.
- Barocas, S. and Nissenbaum, H. (2014). 'Big Data's End Run around Anonymity and Consent', in J. Lane, S. Bender, V. Stodden and H. Nissenbaum (eds.). *Privacy, Big Data, and the Public Good: Frameworks for Engagement* (New York: Cambridge University Press), pp. 44–75.
- Beauchamp, T.L. (2010). 'Autonomy and Consent', in F.G. Miller and A. Wertheimer (eds.). *The Ethics of Consent: Theory and Practice* (New York: Oxford University Press), pp. 55–78.
- Beauchamp, T.L. and Childress, J.F. (2013). *Principles of Biomedical Ethics* (New York: Oxford University Press).
- Davis, S. (2009). 'Is There a Right to Privacy?', *Pacific Philosophical Quarterly* 90 (4): 450–475.
- Fischer, J.M. (2010). 'The Frankfurt Cases: the Moral of the Stories', *The Philosophical Review* 119 (3): 315–336.

- Fischer, J.M. and Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility* (New York: Cambridge University Press).
- Frankfurt, H.G. (1969). 'Alternate Possibilities and Moral Responsibility', in G. Watson (ed.) (2003). *Free Will: Second Edition* (New York: Oxford University Press), pp. 167–176.
- Frankfurt, H.G. (1971). 'Freedom of the Will and the Concept of a Person', *The Journal of Philosophy* 68 (1): 5–20.
- Fried, C. (1968). 'Privacy [A Moral Analysis]', in F.D. Schoeman (ed.) (1984). *Philosophical Dimensions of Privacy: an Anthology* (Cambridge: Cambridge University Press), pp. 203–223.
- Gerstein, R.S. (1978). 'Intimacy and Privacy', *Ethics* 89 (1): 76–81.
- Inness, J. (1992). *Privacy, Intimacy, and Isolation* (New York: Oxford University Press).
- Lever, A. (2013). *On Privacy* (New York: Routledge).
- Macnish, K. (2018). 'Government Surveillance and Why Defining Privacy Matters in a Post-Snowden World', *Journal of Applied Philosophy* 35 (2): 417–432.
- Marmor, A. (2015). 'What Is the Right to Privacy?', *Philosophy & Public Affairs* 43 (1): 3–26.
- McKenna, M. and Pereboom, D. (2016). *Free Will: A Contemporary Introduction* (New York: Taylor & Francis Ltd).
- Mokrosinska, D. (2018). 'Privacy and Autonomy: On Some Misconceptions Concerning the Political Dimensions of Privacy', *Law and Philosophy* 37 (2): 117–143.
- Moore, A.D. (2010). *Privacy Rights: Moral and Legal Foundations* (University Park: The Pennsylvania State University Press).
- Nissenbaum, H. (2009). *Privacy in Context: Technology, Policy, and the Integrity of Social Life* (Stanford, CA: Stanford University Press).
- Parent, W.A. (1983a). 'Privacy, Morality, and the Law', *Philosophy and Public Affairs* 12 (4): 269–288.
- Parent, W.A. (1983b). 'Recent Work on the Concept of Privacy', *American Philosophical Quarterly* 20 (4): 341–355.
- Rachels, J. (1975). 'Why Privacy Is Important', *Philosophy & Public Affairs* 4 (4): 323–333.
- Rickless, S.C. (2007). 'The Right to Privacy Unveiled', *San Diego Law Review* 44 (1): 773–799.
- Roberts, A. (2015). 'A Republican Account of the Value of Privacy', *European Journal of Political Theory* 14 (3): 320–344. <https://doi.org/10.1177/1474885114533262>.
- Rössler, B. (2004). *The Value of Privacy* (Cambridge, UK: Polity Press).
- Rössler, B. (2017). 'Privacy as a Human Right', *Proceedings of the Aristotelian Society* 117 (2): 187–206. <https://doi.org/10.1093/arisoc/aox008>.
- Rössler, B. and Mokrosinska, D. eds. (2015). *Social Dimensions of Privacy: Interdisciplinary Perspectives* (New York: Cambridge University Press).
- Sartorio, C. (2016). 'Frankfurt-Style Examples', in K. Timpe, M. Griffith and N. Levy (eds.). *The Routledge Companion to Free Will* (New York: Routledge), pp. 179–190.
- Solove, D.J. (2008). *Understanding Privacy* (Cambridge, Mass: Harvard University Press).
- Thomson, J.J. (1975). 'The Right to Privacy', *Philosophy & Public Affairs* 4 (4): 295–314.