



A Defense of Privacy as Control

Leonhard Menges¹

Received: 2 March 2020 / Accepted: 24 September 2020 / Published online: 24 October 2020
© The Author(s) 2020

Abstract

Even though the idea that privacy is some kind of control is often presented as the standard view on privacy, there are powerful objections against it. The aim of this paper is to defend the control account of privacy against some particularly pressing challenges by proposing a new way to understand the relevant kind of control. The main thesis is that privacy should be analyzed in terms of source control, a notion that is adopted from discussions about moral responsibility.

Keywords Privacy · Control · Source control

1 Introduction

Many authors have argued that privacy is essentially a certain kind of control. This is the common idea shared by what I will call *control accounts of privacy*. Even though control accounts are often presented as the standard views on privacy, there are powerful objections against them. Most recently, Lundgren (2020) combined two independent and well-known challenges for control accounts into one dilemma: he argues that the best attempt to meet the first challenge is incompatible with the best attempt to meet the second challenge. Moreover, he says that accounts that meet both challenges are not real control accounts.

The aim of this paper is to defend the thesis that privacy is, essentially, a kind of control by proposing a new way to understand the relevant kind of control. Even though my discussion is structured as a reply to Lundgren's paper, the result will be of more general interest because he relies on objections and ideas that have been around for a long time.

Section 2 will present Lundgren's objection. In Sect. 3 I will argue that there is room for the idea that exercising a certain kind of control never diminishes one's privacy. Section 4 is the heart of the paper. In reply to the objection that control

✉ Leonhard Menges
leonhard.menges@sbg.ac.at

¹ Department of Philosophy (KGW), University of Salzburg, Franziskanergasse 1, 5020 Salzburg, Austria

theories are too broad because they imply that privacy is lost in cases in which this is, intuitively, false. I will present and develop what I will call *the source control account of privacy*. In Sect. 5 I will discuss two possible objections against this account. The upshot is that there is a promising and so-far overlooked version of the control account which meets some famous and pressing challenges.

It is important to be clear that this paper is primarily concerned with conceptual, not with axiological or normative issues. From a conceptual perspective we ask what privacy is. From an axiological or normative perspective we can then ask whether privacy, thus understood, is good or bad, whether we have reason to get rid of or support it, and whether we have a right to privacy. The aim of this paper is to explore the concept of privacy. Even though I believe that privacy, understood in the way I will develop below, is, in general, valuable, that we have reason to protect it, and that we have a right to it, these axiological and normative claims are irrelevant for the conceptual aims of this paper (I discuss some valuable aspects of privacy as source control in Menges 2020: 42–45).

A final note is called for. In what follows, I will only argue for the claim that *informational* privacy is essentially a kind of control. Thus, the view I will defend is neutral with regard to the concepts of what is sometimes called locational, bodily, or decisional privacy. The reason for this is pragmatic. It seems sufficiently ambitious to defend control accounts of informational privacy by presenting a new analysis of control. Applying this view to the other notions would require much more space than I have. Thus, I leave these topics for further occasions.

With this in place, let me turn to Lundgren's objection.

2 The Dilemma Objection Against Control Accounts of Privacy

Lundgren's explicit target is the thesis that "*privacy (or the right to privacy) should be conceptualized as some kind of control over some kind of matters*" (Lundgren 2020, 167; italics in the original). He focuses on the concept of privacy (not on the right to it) and begins his objection with a classical argument, developed by William Parent, that I will call the *voluntary divulgence objection*. It says that control accounts are too broad because they imply that a person has privacy even if he "freely divulges everything, no matter how intimate the facts, about himself to a friend" (Parent 1983a, 273; see Lundgren 2020, 168). Parent and Lundgren claim that this implication is implausible and that, therefore, control accounts are implausible. Lundgren says that one can avoid this problem by modifying the main idea of control accounts. The modified version would say that a "person has privacy to the degree that she has control over (access to) her private matters" (Lundgren 2020, 169). But this view will then be in conflict with the best reply to the second objection.

The second objection says that control accounts are too narrow because they imply that our privacy is diminished or lost in situations in which, intuitively, it is not. Lundgren refers to a case presented by Macnish in which "I leave my diary on a table in a coffee shop and return to that shop 30 min later to retrieve it" (Macnish 2018, 420; see Lundgren 2020, 169). I see the diary in the hand of a stranger at

another table. The stranger gives it back to me and convincingly and truly tells me that they have kept it shut. Lundgren and Macnish argue that control accounts imply, implausibly, that by leaving the diary on the table, I have diminished or lost privacy because I have diminished or lost control over (access to) private information. They conclude that control accounts are implausible. As this diary case has striking similarities with cases discussed by Parent that he calls threatened loss cases, I will call Lundgren's line of reasoning the *threatened loss objection* (Davis 2009, 457; Moore 2010, 21 n. 35; Parent 1983b, 344; see also Rickless 2007, 783; the first threatened loss case comes from Thomson 1975, n. 1).

Again, Lundgren contends that a modification of the main idea of control accounts would avoid this objection. The proposal he discusses says that a "person *P* has privacy to the degree that others respect that *P* should have control over (access to) *P*'s private matters" (Lundgren 2020, 169). But this view, he says, comes in conflict with the best reply to the first objection.

For the purposes of this paper, it is not necessary to discuss how Lundgren's proposed solutions to both objections come in conflict with each other. That's because I will argue that Lundgren does not identify the best replies in defense of the claim that privacy should be conceptualized as some kind of control. That is, the best reply to the first objection is, *pace* Lundgren, compatible with the best reply to the second objection. But before showing this, let me briefly come to the last part of Lundgren's objection because it will also be relevant for the remainder of this paper.

Lundgren argues that there are accounts of privacy that look like control accounts and that avoid the dilemma. He presents one which has some similarities with Andrei Marmor's (2015) recent view. It says that a "person has privacy to the degree that she has control over how she can present herself to others" (Lundgren 2020, 172). Lundgren then argues that this and similar views are not, in fact, control accounts, but versions of limited access accounts. How a person can present herself to others, the argument goes, completely depends on what information they have already accessed, not on how much control the person has. Thus, whether or not a person has privacy "does not seem to be affected by less or more control, but by less or more limited access" (Lundgren 2020, 172). Therefore, the objection goes, the account is not a real control account. Rather, it collapses into a limited access account, which is typically thought of as the major opponent of control accounts. I will call this line of thinking the *collapse objection*.

Lundgren concludes that the dilemma "can only be avoided by giving up the concept of control in favor of limited access" (Lundgren 2020, 173). His main thesis is that privacy should not and cannot be defined in terms of control (see also 2020, 173) or, in other words, that it is false to claim that having privacy is having a kind of control. In what follows, I will defend the main idea of control accounts that privacy is a kind of control by presenting a new way to understand the relevant notion of control. The resulting view will avoid the dilemma without collapsing into limited access accounts.

3 The Voluntary Divulgence Objection

In this section I will flesh out three voluntary divulgence cases and I will argue that none of them poses a problem for the idea that having privacy is, essentially, having a kind of control. Then, I will suggest how to generalize this conclusion to all similar cases. But let me begin with Lundgren's version of the voluntary divulgence objection:

Parent's counterexample was that according to some control accounts of privacy (i.e., those which equate privacy with control) I cannot diminish my privacy by sharing private matters, if I am in control when doing so. This is counterintuitive, since I obviously can diminish my privacy by sharing private matters (this needs to be recognized by any proponent of a control account—and most do—since it reduces my control of my private matters) (Lundgren 2020, 171).

Based on this line of reasoning, Lundgren concludes that the main idea of control accounts needs to be modified.

A point of agreement is that every account of privacy should make room for its being possible to diminish one's own privacy. Consider—and this is the first case—a truth drug that makes people want to tell others about intimate facts. A person who divulges personal information to others because of the drug plausibly diminishes her own privacy. Control accounts have no problem with such a case. According to them, privacy is a kind of control over (access to) information. When people tell others about personal information because of truth drugs, they do not exercise the relevant kind of control. The truth drugs diminish their control over how personal information flows and, thereby, their privacy is diminished. The same holds for people who reveal everything about their lives because of mental problems, ignorance, alcohol, or other control-diminishing factors. Thus, in general, control accounts have no problems making sense of its being possible to diminish one's own privacy.

Proponents of the voluntary divulgence objection, however, consider cases in which people exercise the kind of control that, according to control accounts, constitutes their privacy. Then they claim that, intuitively, this diminishes their privacy. To illustrate, consider the following second voluntary divulgence case: After carefully considering the situation, I freely and willingly tell my friend that I am going to become a parent, that I want to marry my partner, and that I have certain health problems.

Let us assume that I have full control over (access to) the information about my family planning, romantic love, and health issues, and that I exercise this control by what one may call voluntarily divulging personal information about myself to a friend. Parent seems to have cases of this kind in mind when he imagines a person who "freely divulges everything, no matter how intimate the facts, about himself to a friend" (Parent 1983a, 273). But do I, thereby, diminish or lose privacy? This is the decisive question. The voluntary divulgence objection has bite if the most intuitive answer is "Yes". But I find this far from obvious. And I'm

not the only one. For example, Julie Inness says about similar cases that “we are including another within our realm of privacy, not lessening our privacy” (Inness 1992, 46).¹

Of course, I diminish or lose some things in these cases. For example, I diminish the range of personal information that is kept secret from others. I also lose my ability to make sure whether or not others learn about certain information, and we may call this ability some kind of control (I will come back to this kind of control in the next section). But it is far from clear that losing these things involves losing privacy.²

Let us now consider cases in which people divulge intimate facts to complete strangers, for example on a website. Do they diminish privacy? The answer I find most plausible is this: if they exercise the same kind of control that I exercise when I tell my friend that I am going to become a parent, want to marry my partner, and have certain health problems, then they do not diminish their privacy.

At first sight, this may look counter-intuitive. It seems plausible that those who present their entire life on the web diminish their privacy. Note, however, that the view I am arguing for does not need to deny this. In many cases in which people present their intimate thoughts, feelings, history, and so on to strangers, there is reason to doubt that they have full control over what they do. In some cases, people have mental health issues and in many cases they do not fully understand how the Internet or the platform they are using works. These are control-diminishing factors. Then, they do not exercise the same kind of control that I exercise when I tell my

¹ An anonymous referee objects that Inness' and my reply rests on the assumption that privacy is valuable. This would come into conflict with my announcement in the Introduction that axiological claims are irrelevant for my main arguments. If it is correct that my reply rests on value claims, then I would be willing to give up the aim to stay neutral towards axiological questions. However, I do not see that the reply rests on assumptions about the value of privacy. We can distinguish between the intuition that privacy has been retained in this situation and the intuition that something valuable has been retained. Personally, I have both. However, the value intuition (that something good has been retained) is not necessary for the point which is most relevant in this context. The important point is that privacy was retained when I told my friend about personal information. It seems possible that one can have this intuition even if one does not think that something good has been retained.

² Here is a possible objection against my claim that I do not lose privacy by voluntarily divulging the relevant information to my friend: (1) One can exercise full control over (access to) information by sharing privacy with others and I do this by voluntarily divulging the information to my friend; (2) sharing privacy with others involves diminishing privacy. Therefore, one can exercise full control over (access to) information and, thereby, diminish one's privacy in some voluntary divulgence cases. As a reply, both premises are not sufficiently clear or convincing. Regarding (1), it is not clear how, exactly, to make sense of the sharing metaphor with regard to privacy: assuming that privacy is a state one can be in (e.g., others having restricted access) or an agential property (e.g., one's having control), it is unclear how privacy itself can be shared. Regarding (2), whether sharing privacy involves diminishing privacy depends on the nature of privacy. If privacy is more like intimacy, then we can share it without losing it (see, e.g., Gerstein 1978; Inness 1992). If privacy is more like secrecy, then sharing it, plausibly, involves losing it. At this point it is unclear whether privacy is more like intimacy or like secrecy or like both. Thus, at this point, (1) and (2) are not sufficiently clear or convincing. Therefore, I take the claim to be defended that my voluntarily divulging personal information to my friend does not diminish my privacy. Thanks to a referee for pressing me on this.

friend about intimate facts. Therefore, the view I argue for can agree that these people diminish their privacy.

Now, consider—as a third voluntary divulgence case—a person who has complete control and exercises it by revealing intimate facts to the public. To be realistic, take Peter Railton’s admirable Dewey Lecture (2015). In the lecture, Railton presents a series of moments from his personal life that constitute “a transition from insider to outsider, or back” (2015, 2). The final transition is constituted by his giving this very talk and then allowing others to upload the manuscript. That’s because he talks openly about his depression and, in particular, his “fear of social embarrassment and humiliation” (2015, 13). He says: “I now give to all of you my experience, as story, a tale, an example, you might tell others, or yourself, in order to open a non-threatening conversation with yourself or others about what seeking help can do” (Railton 2015, 15).

Let us assume that, when giving the lecture and allowing others to upload the manuscript, Railton had full control over what he did. Did he, thereby, diminish his privacy? Again, I find it far from obvious that the most intuitive answer is “Yes”. Of course, he diminished the range of personal information about him that is kept secret from others and lost the ability to make sure whether or not others learn about his personal life. But it is, as I said above, not clear that losing these things involves losing privacy. Again, it also seems sensible to say that Railton included the philosophical community within his private realm or that he shared private information with the community in order to, among other things, sensitize us to mental health problems.

At this point, opponents of control accounts may agree that neither Railton nor I diminish privacy in the two cases by exercising full control over (access to) the relevant information. They may contend, however, that the voluntary divulgence objection works as soon as there is *one* case in which agents exercise full control over (access to) personal information, share the information with others and, thereby, diminish their privacy. The objection concludes that I have not shown that there is no case of this kind (thanks to an anonymous referee for this suggestion).

Opponents of the control account would be right about this, and I am not able to show that there is no case of this sort—impossibility proofs are notoriously hard in philosophy. But let me present two weaker replies. First, the burden of proof seems to lie on the side of the opponent of control accounts now. Parent and Lundgren sketch the form of a voluntary divulgence case. I have fleshed out three versions of it (the person on truth drugs, my telling a friend about personal issues, and Railton’s Dewey Lecture) and I have argued that none of them poses a problem for the control account. Now, it is my opponent’s job to show that there are voluntary divulgence cases that work against the control account.

Second, what I have said so far suggests a general way to deal with voluntary divulgence cases. Proponents of the control account should argue in the following way. (1) In the Dewey Lecture case: Railton’s exercising the relevant kind of control over personal information does not diminish his privacy. (2) For any voluntary divulgence case: if the agents exercise the same kind of control that Railton exercises by giving the lecture, then there is no privacy-relevant difference between this case and the Dewey Lecture case. Therefore, for any voluntary divulgence case: if

the agents exercise the same kind of control that Railton exercises in the Dewey Lecture case, then the agents do not diminish their privacy.

The success of this argument by analogy rests on the plausibility of premise (2). In order to test it, we need more voluntary divulgence cases. Providing them is, as I said in the first reply, the job of opponents of the control account. In the meantime, this argument looks promising.

With this in mind, let me turn to the next objection that will give me the opportunity to present the promised more substantial account of control.

4 The Threatened Loss Objection

The second horn of Lundgren's dilemma is a threatened loss objection. It relies on cases in which personal information about me, for example my diary, is readily available to a stranger, who knows that this is so. The objection says that this fact diminishes my control over (access to) the information and that control accounts imply that this diminishes my privacy. However, as long as the stranger does not access the information—say, as long as they do not read my diary—this is implausible. There is a threat of privacy loss in these cases, but no actual loss. The objection concludes that control accounts are implausible.

My main reply to the threatened loss objection is that control accounts are not committed to the claim that privacy is diminished in cases of this kind. I present the basic idea in the following sub-section (I discuss threatened loss cases and some aspects of the source control account at length in Menges 2020). In Sect. 4.2 I apply the resulting view to several cases, including voluntary divulgence cases. In Sect. 4.3 I discuss the relation between the concepts of privacy, responsibility, and autonomy, and in Sect. 4.4 I argue that the view developed here is a new one.

4.1 The Source Control Account of Privacy

Generally, there are different ways to spell out the notion of control. The threatened loss objection is only successful for a very specific understanding of control, namely one that involves having effective voluntary choice over whether or not others learn about the information. Having effective voluntary choice is the kind of control we have in mind when we say that people have control over whether they become a lawyer or dentist, eat vanilla or chocolate ice cream, accept or reject an offer. It is very plausible that I lose effective voluntary choice over whether or not the stranger learns about the information when I leave the diary in the café. This is so because there is nothing I can do to prevent the stranger from reading the diary. However, there is no need for control theorists to account for privacy in terms of having effective voluntary choice. They can and should opt for another kind of control.

Let me present considerations about control over what we do as a model for control accounts of privacy. Consider the following case first developed by Locke (1706, bk. II, chap. 21):

I find myself in a room with a person I very much like to be with and, therefore, I freely decide to stay there. Unbeknownst to me, however, the only door out of the room is locked.

Intuitively, I exercise an important kind of control over what I do by staying in the room. I am not forced to do it and I do not experience a compulsion to stay. I freely decide to do it. However, and unbeknownst to me, I never had effective choice over leaving or staying because the only door was locked. Thus, I exercise an important kind of control by staying without having effective voluntary choice over whether or not I stay.

Now take a structurally similar case:

Jones has resolved to shoot Smith. Black has learned of Jones's plan and wants Jones to shoot Smith. But Black would prefer that Jones shoot Smith on his own. However, concerned that Jones might waver in his resolve to shoot Smith, Black secretly arranges things so that, if Jones should show any sign at all that he will not shoot Smith (something Black has the resources to detect), Black will be able to manipulate Jones in such a way that Jones will shoot Smith. As things transpire, Jones follows through with his plans and shoots Smith for his own reasons. No one else in any way threatened or coerced Jones, offered Jones a bribe, or even suggested that he shoot Smith. Jones shot Smith under his own steam. Black never intervened (McKenna and Coates 2020, sect. 3.2).

This is a so-called Frankfurt case that is well known from debates about free will and responsibility (for the original, see Frankfurt 1969). As Jones follows through with his own plan and shoots Smith because he wants to do it and decides to do it without any intervention, he exercises an important kind of control when he shoots Smith. However, he never had effective choice over whether or not he shoots Smith because of Black. Thus, we can have an important kind of control over what we do without having effective choice over whether or not we do it.

Let me call the kind of control that I have in the closed-door case and that Jones has in the Frankfurt case *source control* (McKenna and Pereboom 2016, 39 call it "source freedom"). The basic idea is that agents have source control over what they do just in case they are the right kind of source of their doing it. What they do must be grounded, *in the right way*, in them, and not in someone or something else. Importantly, for having responsibility-relevant source control over an action, it is not sufficient that the person is the source of the action. This is because a person can be the source of an action in a variety of ways. When brainwashed agents kill someone else because they were brainwashed, then they are still *somehow* the source of the killing. However, they are not the source of the killing *in the way* that makes them responsible for the killing. Thus, the expression "the right kind" in "an agent's being the right kind of source of an action" is important. It is meant to be a place holder for the kind of source that is necessary for the person's being responsible for their actions.

When, exactly, is an agent the right kind of source of an action in order to be considered responsible for it? This is a complicated and hotly debated matter (for an overview, see McKenna and Coates 2020). Prominent suggestions are that

actions must be grounded in certain desires or cares of the agents (e.g., Arpaly and Schroeder 2014), in their beliefs about what is valuable (e.g., Watson 1975), or in some specific dispositions or traits, often called reason-responsiveness (e.g., Fischer and Ravizza 1998). A famous view says, for example, that the action must be grounded in the agents' having the first order desire to perform it and their having the second order desire that they have the first order desire to perform the action (see Frankfurt 1971). This is, of course, not the place to delve into this debate. But the important point for my purposes is that one can have source control over what one does without having effective voluntary choice over whether or not one does it.

My main proposal is that privacy theorists can and should spell out privacy in terms of source control. According to the resulting *source control account of privacy*, an agent has privacy with regard to a certain piece of information just in case the person is the right kind of source of the relevant information flow if the information flows at all. In other words: an agent's having privacy with regard to a piece of information consists in the agent's being such that if the information flows to others, then the agent is the right kind of source of this information flow.³

When, exactly, is an agent the right kind of source of an information flow for its being true that the agent has privacy with regard to the relevant information? Must the information flow be grounded in some of the agent's desires, cares, or beliefs, the agent's reason-responsiveness, quality of will, or something completely different? A full-blown version of the source control account of privacy needs to answer this question. Admittedly, I do not know the answer, and if I knew it, developing it would, plausibly, require several papers (compare the corresponding debate between responsibility scholars).

Even though this is not completely satisfying, the source control account is still a big step forward. This is because, first, privacy scholars typically do not say much at all about how to understand the relevant kind of control (see Sect. 4.4). Therefore, arguing that the relevant kind of control should be understood in terms of source control is much more specific than most privacy scholars have been so far. Second, the main aim of this paper is to defend the control account against pressing objections. The defense says that these objections lose their force if one understand privacy in terms of source control. This defense works regardless of whether source control is spelled out in terms of, say, reason-responsiveness, cares, or beliefs.

In what follows, I will illustrate the source control account of privacy by spelling out source control in terms of first- and second-order desires. Thus, I will assume for illustration, that we have source control over a piece of information just in case if that information flows to others at all then this flow is grounded in our first order desire that it flows in this way and in our having the second order desire that we have that very first order desire. It is important to keep in mind, however, that *this is not the view I argue for*. I argue for the more general thesis that having privacy

³ Note that we can *have* privacy with regard to a piece of information at a certain point in time even if we do not *exercise* control. When I am asleep, for example, I do not exercise control over information. But as long as the piece of information does not flow to anyone else, then I still have privacy with regard to it, according to the source control account.

is essentially having source control over information flow. The Frankfurt-inspired idea to spell this out in terms of first- and second-order desires is meant to be a mere illustration of how the general thesis *could* be specified.

4.2 Some Cases

Let me illustrate the source control account by applying it to some cases. Consider, first, the diary case which is a classical threatened loss case. As I leave the diary in the café, I have no effective choice over whether or not other people learn about the information in it. Thus, control as effective choice is lost. However, as nobody learns about the information, the information does not flow to anyone. I can still be the first who tells the stranger who holds my diary everything that is written in it. Thereby, I can still be the right kind of source of the information flow. Thus, according to the source control account, my privacy has not been diminished in the diary case. This is the intuitively correct result, which suggests that standard threatened loss cases pose no problem for the source control account of privacy.

Second, imagine that “a couple accidentally reveals intimate facts about their intimate relations (perhaps by accidentally leaving a live-streaming web camera on)” (Lundgren 2020, 171). As the revelation is *accidental*, the flow of information is, plausibly, not grounded in what they desire and in what they desire to desire. Thus, the couple is not the right kind of source of the information flow and, thereby, their privacy is diminished, according to this specific version of the source control account. Again, this is the intuitively correct result.

Third, consider a variation of the case: the couple desires to reveal intimate information about their relationship by turning the camera on and they desire to have this desire. Then, they forget about this plan and do not turn the camera on. However, a bug in their computer program turns on the camera such that intimate information about them is revealed. As the flow of information is not grounded in the couple in the right way—that is, the information does not flow *because of* their desires—the source control account says that their privacy has been diminished.

I have learned that people have different intuitions about this case, but I find this conclusion plausible. A possible source of resistance may be the idea that claiming that privacy is lost commits one to saying that something valuable has been lost or that a right has been infringed. And this does not seem to be the case. However, these claims should be kept apart. The source control account only says that privacy is diminished in the case at hand. It neither follows that the couple has lost something valuable nor that their rights have been infringed or violated. Whether this is so is an independent question.⁴

⁴ Consider the following analogy: I own a photo of my ex-partner, I desire to throw it away and I have the desire to have this desire. However, I forget to throw it away and then accidentally lose it while moving apartments. Intuitively, I lost private property, but nothing bad happened and no right was infringed. This case is analogous to the one of the couple who accidentally lost privacy even though they wanted to reveal the relevant information.

Let me finally turn to voluntary divulgence cases again. Recall my telling my friend that I am going to become a parent, that I want to marry my partner, and that I have certain health problems. And recall Railton's Dewey Lecture in which he reveals information about his personal history. The source control account says that as long as I am the right kind of source of sharing these things—for example, I desire to share them and I desire to have that desire—I do not diminish my privacy by revealing these intimate facts. Similarly, it says that Peter Railton did not diminish his privacy when he exercised source control by giving the Dewey Lecture. Here again, the source control account has the intuitively plausible results.

Now, we can also spell out the general reply to voluntary divulgence cases in more detail. (1) In the Dewey Lecture case: Railton's exercising source control over personal information does not diminish his privacy. (2) For any voluntary divulgence case: if the agents exercise the same source control that Railton exercises by giving the lecture, then there is no privacy-relevant difference between this case and the Dewey Lecture case. Therefore, for any voluntary divulgence case: if the agents exercise the same source control that Railton exercises in the Dewey Lecture case, then the agents do not diminish their privacy.⁵

To sum up, in the cases discussed here, the source control account has the intuitively correct implication. In the threatened loss and in the voluntary divulgence cases privacy is not diminished, in the accidental revelation cases it is.

4.3 Privacy, Responsibility, and Autonomy

The source control account of privacy is inspired by source control accounts of responsibility. Moreover, source control accounts of responsibility in general and the specific Frankfurt-inspired version that I have sketched above that spells out control in terms of first- and second-order desires are often discussed as accounts of personal autonomy. This view says, roughly, that acting autonomously just is acting from the relevant first- and second-order desires (see the essays in Buss and Overton 2002). These observations raise the question of how privacy, responsibility, and autonomy relate to each other.

Generally, how the three notions relate to each other depends on how, exactly, source control is spelled out in the context of privacy, responsibility, and personal autonomy. It may turn out that the three notions come very close together. This would be so if the independently best account of privacy yields that privacy is to be analyzed in terms of a certain kind of source control and if the independently best accounts of responsibility and personal autonomy yield that these notions should be spelled out in terms of exactly the same kind of source control. Then,

⁵ Note, however, that the source control account leaves room for its being possible that one fails to protect one's privacy (thanks to an anonymous referee for pressing me on this). Imagine that I desire and desire to desire that Snape does not learn what I'm doing tonight and I try to fulfill that desire by closing the curtain. However, I fail because, unbeknownst to me, there is a hole in the curtain and Snape looks through it. The Frankfurt-inspired version of the source control view would say that my privacy has been diminished because the information flows without my being the right kind of source of the flow of information about what I am doing tonight..

the conditions for autonomous action would be identical with the control condition of responsibility and with the kind of control that constitutes having informational privacy.

Even if it is possible that the three notions are so similar, it seems unlikely to me that personal autonomy and responsibility on the one hand come that close to privacy on the other. Before showing why, let me briefly comment on the relation between personal autonomy and responsibility. Most accounts of the control condition of responsibility can also be interpreted as accounts of autonomous agency (see Buss and Westlund 2018 for an overview). There are surely interesting differences between the two, but diving into them would lead us too far away from the main topic (but see, e.g., Wallace 1994, chap. 3). Therefore, I will assume that autonomous agency just is the kind of control that is necessary for being responsible.

Responsibility for actions is, on a widely held view, essentially concerned with when it is appropriate to respond to the agents' performing bad actions in a negative, unwelcome way, for example by blaming or sanctioning the agents (see, e.g., the classical Strawson 1962). Part of what explains why agents need to have a strong kind of control for being responsible is that it would be unjust, unfair, or undeserved to blame or sanction a person for an action if she did not have a strong kind of control over performing it (see, e.g., Wallace 1994; Strawson 1994; Watson 1996; Pereboom 2014). This is why the first- and second-order desire account of source control is not very popular among recent responsibility scholars: these desires do not provide the agents with enough control over what they do for its being fair, just, or deserved to blame or sanction them (see, e.g., Wallace 1994, chap. 3; Mele 2019, chap. 2).

While considerations about when it is fair, just, or deserved to respond in unwelcome ways to what agents do explain why responsibility requires a strong kind of control, there is no analogous reason to think that having privacy requires a similarly strong kind of control. This becomes most vivid when we imagine that our universe is deterministic. Some doubt that and it is reasonable to discuss whether humans can be responsible in a deterministic universe. The main worry is that determinism may rule out the kind of control that is necessary for fair, just, or deserved blame (for overviews see Clarke and Capes 2017; Caruso 2018). However, nobody doubts and it does not seem reasonable to discuss whether humans can have informational privacy in a deterministic universe. This suggests that the kind of control that constitutes having privacy is weaker than the kind of control that is necessary for responsibility and autonomous agency. If this is so, then the notions of responsibility and autonomy on the one hand and the notion of privacy on the other are not very close to each other.

To sum up, source control plays important roles in different philosophical debates, most prominently about responsibility and personal autonomy. Even though the source control account of privacy is inspired by these discussions, there is reason to think that the best way to spell out the kind of source control that is relevant for privacy differs in important respects from the kind of source control that is relevant for responsibility and autonomy.

4.4 Is the Source Control Account of Privacy Novel?

So far, I have presented the source control account of privacy as a specific version of the thesis that privacy is essentially a kind of control. Moreover, the source control view avoids important objections that have been raised for traditional control accounts. Now, one may ask how, exactly, “traditional” control accounts differ from the source control view: is the source control account really a new one or is it what proponents of control accounts had in mind anyway? In what follows, I will discuss three representative control accounts and then explore a line of thinking that suggests that no account has, so far, made sense of privacy in terms of source control.

Many control theories do not make it explicit whether they spell out control in terms of effective voluntary choice or in terms of being the right kind of source. Beate Rössler, for example, who is one of the most prominent recent control theorists, proposes “the following definition of privacy: something counts as private if one can oneself control the access to this ‘something’” (Rössler 2004, 8). She then demonstrates this general view: “[t]ake data or information about me: here too it is reasonable to say that this is private if I can and/or should control access to it” (Rössler 2004, 8). Rössler goes on to distinguish informational, local, and decisional privacy and to discuss whether the concept of privacy is inherently normative. But she does not discuss whether the relevant kind of control should be thought of as being able to effectively choose whether or not something happens or as being the right kind of source of an event. Thus, it is hard to say whether she defines having (informational) privacy in terms of having effective voluntary choice or in terms of being the right kind of source of information flows.

Some proponents of control accounts are more explicit. Most clearly, Charles Fried contends that “[p]rivacy, thus, is control over knowledge about oneself”, and he specifies that “[t]he person who enjoys privacy is able to grant or deny access to others” (Fried 1968, 210). Being able to grant or deny access is a form of control as effective voluntary choice. Thus, Fried makes sense of having privacy in terms of having effective voluntary choice, not in terms of being the right kind of source.

Some of those who focus on the importance of the right to privacy also suggest that effective voluntary choice is essential for this right. Marmor, for example, argues that the right to privacy is important because it protects our ability to shape different kinds of relationships (see also Rachels 1975). Marmor contends that

everyone needs some choice about *how close or how distant they want to be* from different others. A reasonable amount of control over ways in which we present ourselves to others is necessary for the kind of choices we want to make about the social interactions we have with different people (Marmor 2015, 11, italics added).

Marmor suggests that the right to privacy protects “choice about how close or how distant” others are from us. This is effective voluntary choice. Thus, even if Marmor does not aim at an analysis of the concept of privacy, he seems to think that the right to privacy is concerned with effective voluntary choice.

As an intermediate conclusion, many proponents of control accounts do not say anything about whether they understand control in terms of voluntary choice or

sourcehood. Those who are more explicit tend to present control as effective voluntary choice. If this is on the right track, then the source control view, really is a novel version of the control account of privacy.

Moreover, there is indirect evidence for the claim that most proponents of the control view implicitly understand privacy as effective voluntary choice. The idea is this: if one holds the view that having privacy is to be spelled out in terms of having source control, then it is very easy to reply to the threatened loss objection that has been around for more than 40 years. The response is that privacy is not diminished in these cases because source control is not diminished in them. As I am not aware of any author who has proposed this response to the threatened loss objection, it seems plausible to conclude that no author has yet thought of privacy in terms of source control.

Let me briefly summarize the most important results so far. Proponents of control accounts of privacy can analyze having privacy in terms of having effective voluntary choice or in terms of being the right kind of source of information flows. Opting for source control circumvents the classical voluntary divulgence and the threatened loss objection. The view is inspired by accounts of moral responsibility and personal autonomy, but there is independent reason to think that these three notions do not fall together. Moreover, there is good reason to think that the source control view has so far been overlooked in the conceptual landscape.

5 Further Objections

5.1 The Collapse Objection

Recall the structure of Lundgren's dilemma objection: the best reply to the voluntary divulgence objection involves revising the main idea of control accounts; the best reply to the threatened loss objection involves another revision of the main idea; however, both revisions come in conflict with each other. So far, I have argued that the best replies to both objections do not require revising the core idea that having privacy is, essentially, having a kind of control. Thus, there is no dilemma for control accounts. However, Lundgren presents a third objection. The general idea is that accounts that avoid the dilemma are not real control accounts but collapse into other views on privacy.

Note that Lundgren presents the collapse objection as a problem for a version of the control account that is inspired by Marmor's recent view on the right to privacy. But he suggests that the objection can be generalized for all control accounts that avoid the voluntary divulgence and the threatened loss objection (see Lundgren 2020, n. 17). The following is meant to test if the collapse objection hits the source control account. In order to do so, I will slightly reformulate the objection:

However, [the source control account] also changes the conception of control as such. My [being the right kind of source of an information flow] is not affected by the fact that someone could interfere with it (as the concept of control discussed by Macnish would). It is affected only by actual interfer-

ence. More importantly, it is affected by access alone. That is, the [source control account] of privacy does not seem to be affected by less or more control, but by less or more limited access, since privacy—in the given example—is affected by another person’s access alone. Thus, the [source control account] of privacy resolves the dilemma by taking the form of a limited access account (Lundgren 2020, 172).

First, the kind of control that Lundgren has in mind is, most probably, very different from the kind of control the source control account refers to. But this is not an objection against the source control account. It only shows that proponents and opponents of control accounts have, so far, not considered an attractive way to spell out the notion of control.

Second, and very generally, every account of privacy should say that accessing personal information sometimes diminishes a person’s privacy. If the stranger reads my diary before I come back and, thereby, accesses personal information about me, she obviously diminishes my privacy. Every account should make sense of this. Control theorists define privacy in terms of control. Therefore, they can and, I believe, should say that accessing information diminishes privacy only if it diminishes a certain kind of control.

More specifically, the source control account says that people’s privacy is diminished when personal information about them flows to other agents even though they were not the right kind of source of this flow. A stranger’s accessing the information is a paradigmatic way of letting information flow in such a way. According to the source control account, the stranger’s accessing personal information about me by reading my diary that I forgot in the coffee shop diminishes my privacy because it affects my source control over personal information. Thus, on this view, access affects privacy only if it affects source control.

This shows that the source control account does not collapse into the limited access account. A central implication of the former is that diminishing source control necessarily diminishes privacy and the other way around. As it is possible to diminish source control by accessing information, accessing information can diminish privacy. But it only does so *by* diminishing source control. Therefore, the most fundamental privacy-relevant aspect is, according to this view, source control and not access alone.

5.2 The No Control Objection

Even if one grants the view presented here that it deals nicely with some classical objections and that it is novel and distinct from other accounts, one may still object that it is misleading to call it a *control* account of privacy. It may be sensible to analyze having privacy in terms of being the right kind of source of information flows, the objection goes, but it is not sensible to call this a version of the control account. If this objection holds, then everything that has been said so far does not amount to a defense of the main thesis that privacy is essentially a kind of control.

As a preliminary reply, one *may* accept the source view presented here and agree that we should not call it a version of the control view. The most important

conceptual task is to find out what privacy is. The account presented here says that one's having privacy is, essentially, one's being the right kind of source of information flows if the information flows at all. This is an interesting, innovative, and plausible claim even if it is not put in terms of control.

This being said, there is good reason to call sourcehood a kind of control. To see this, recall the Frankfurt and the closed-door case presented in Sect. 4.1. In these cases, it seems as if the agents cannot effectively choose between different options. I cannot choose whether or not I leave the room, Jones has no effective choice between shooting or not shooting Smith. However, it is very intuitive that I exercise an important kind of control by staying in the room and that Jones exercises an important kind of control by shooting Smith. The kind of control we exercise can be called source control: we have some kind of control over what we do by being the right kind of source of what we do.

The source control account of privacy says that what constitutes having privacy is analogous to my being the source of staying in the room and to Jones' being the source of shooting Smith. As it is plausible to call the kind of sourcehood which is relevant in the latter cases "control" it is also plausible to call the kind of sourcehood which is relevant for privacy "control".

Moreover, many debates about control in the realm of responsibility are concerned with sourcehood. One of the most influential and hotly debated books on responsibility is John Martin Fischer's and Mark Ravizza's *Responsibility and Control: A Theory of Responsibility* (1998). They argue, roughly, that the kind of control that partly grounds moral responsibility for actions consists in the action's being caused by a reason-responsive mechanism of the agent. If the action is caused by such a mechanism then the agent is the right kind of source of the action that may ground the agent's being responsible for it. Thus, Fischer and Ravizza analyze control in the realm of responsibility in terms of the sources of actions. The view on privacy presented here analyzes privacy in terms of the sources of information flows. As many people call sourcehood in the realm of responsibility a kind of control, it also makes sense to call sourcehood in the realm of privacy a kind of control.

More generally, it is possible to analyze privacy in terms of sourcehood without using the notion of control. However, there is good reason to use this notion and to call the resulting view the *source control account of privacy*.

6 Conclusion

The aim of this paper is to defend control accounts of privacy against some important objections that have recently regained prominence, thanks to Lundgren's "A Dilemma for Privacy as Control". I have tried to achieve this goal by presenting and defending the main idea of a new version of the control account, namely the source control view.

Let me close the paper by pointing out further research questions, some of which I have already indicated. First, I have not discussed the question of whether privacy is a neutral, value-laden, or normative concept. This should be explored from the perspective of the source control account. Second, we need to spell out how, exactly,

to understand sourcehood in the realm of privacy. Is it a matter of the agent's having certain cares, beliefs, or desires, of the agent's being reason-responsive, or something else entirely? Third, we need to know how informational privacy, thus understood, relates to what is sometimes called decisional and locational privacy. Fourth, from an axiological and normative perspective, we should explore how valuable privacy, thus understood, is and whether we have a right to it. Fifth, the resulting view should help make sense of specific problems in the context of, for example, information technologies, medicine, and politics. Finally, it should be explored whether this view is superior to other important accounts of privacy such as the limited access view or those accounts that avoid classical analyses of the concept of privacy (see, e.g., Solove 2008; Nissenbaum 2009; Allen 2011). Thus, much needs to be done.⁶

Funding Open access funding provided by Paris Lodron University of Salzburg.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Allen, Anita. 2011. *Unpopular privacy: What must we hide?* New York: Oxford University Press.
- Arpaly, Nomy, and Timothy Schroeder. 2014. *In praise of desire*. New York: Oxford University Press.
- Buss, Sarah, and Lee Overton (eds.). 2002. *Contours of agency: Essays on themes from Harry Frankfurt*. Cambridge, Mass.: MIT Press.
- Buss, Sarah, and Andrea Westlund. 2018. Personal autonomy. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2018. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2018/entries/personal-autonomy/>.
- Caruso, Gregg D. 2018. Skepticism about moral responsibility. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2018. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2018/entries/skepticism-moral-responsibility/>.
- Clarke, Randolph, and Justin Capes. 2017. Incompatibilist (Nondeterministic) theories of free will. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2017. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2017/entries/incompatibilism-theories/>.
- Davis, Steven. 2009. Is there a right to privacy? *Pacific Philosophical Quarterly* 90 (4): 450–475.
- Fischer, John Martin, and Mark Ravizza. 1998. *Responsibility and control: A theory of moral responsibility*. New York: Cambridge University Press.
- Frankfurt, Harry G. 1969. Alternate possibilities and moral responsibility. In *Free Will*, edited by Gary Watson, 2nd ed., 167–76. New York: Oxford University Press, 2003.
- Frankfurt, Harry G. 1971. Freedom of the will and the concept of a Person. *The Journal of Philosophy* 68 (1): 5–20. <https://doi.org/10.2307/2024717>.

⁶ I am grateful for helpful comments from Peter Königs and Björn Lundgren. Special thanks to two anonymous referees for this journal for feedback that helped make the paper much better. As always, thanks to Claire Davis for proofreading.

- Fried, Charles. 1968. Privacy [A moral analysis]. In *Philosophical dimensions of privacy: An anthology*, edited by Ferdinand David Schoeman, 203–23. Cambridge: Cambridge University Press, 1984.
- Gerstein, Robert S. 1978. Intimacy and privacy. *Ethics* 89 (1): 76–81.
- Inness, Julie. 1992. *Privacy, intimacy, and isolation*. New York: Oxford University Press.
- Locke, John. 1706. *An essay concerning human understanding*. Edited by Peter H. Nidditch. Oxford: Clarendon Press, 1982.
- Lundgren, Björn. 2020. A dilemma for privacy as control. *The Journal of Ethics* 24 (2): 165–175. <https://doi.org/10.1007/s10892-019-09316-z>.
- Macnish, Kevin. 2018. Government surveillance and why defining privacy matters in a post-snowden World. *Journal of Applied Philosophy* 35 (2): 417–432. <https://doi.org/10.1111/japp.12219>.
- Marmor, Andrei. 2015. What is the right to privacy? *Philosophy & Public Affairs* 43 (1): 3–26.
- McKenna, Michael, and D. Justin Coates. 2020. Compatibilism. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2020. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/spr2020/entries/compatibilism/>.
- McKenna, Michael, and Derk Pereboom. 2016. *Free will: A contemporary introduction*. New York: Taylor & Francis Ltd.
- Mele, Alfred R. 2019. *Manipulated agents: A window to moral responsibility*. New York: Oxford University Press.
- Menges, Leonhard. 2020. Did the NSA and GCHQ diminish our privacy? What the control account should say. *Moral Philosophy and Politics* 7 (1): 29–48. <https://doi.org/10.1515/mopp-2019-0063>.
- Moore, Adam. 2010. *Privacy rights: Moral and legal foundations*. University Park: The Pennsylvania State University Press.
- Nissenbaum, Helen. 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 & Public Affairs* 12 (4): 269–288.
- Parent, W.A. 1983b. Recent work on the concept of privacy. *American Philosophical Quarterly* 20 (4): 341–355.
- Pereboom, Derk. 2014. *Free will, agency, and meaning in life*. New York: Oxford University Press.
- Rachels, James. 1975. Why privacy is important. *Philosophy & Public Affairs* 4 (4): 323–333.
- Railton, Peter. 2015. Innocent abroad: Rupture, liberation, and solidarity (Dewey Lecture). <https://www.dropbox.com/s/1vuvmyohndrsy/Railton%20Dewey%20Lecture%20Central%20APA%202015%20revised.pdf?dl=0>.
- Rickless, Samuel C. 2007. The right to privacy unveiled. *San Diego Law Review* 44 (1): 773–799.
- Rössler, Beate. 2004. *The value of privacy*. Cambridge, UK: Polity Press.
- Solove, Daniel J. 2008. *Understanding privacy*. Harvard University Press.
- Strawson, Galen. 1994. The impossibility of moral responsibility. In *Free Will*, edited by Gary Watson, 212–28. New York: Oxford University Press, 2003.
- Strawson, Peter F. 1962. Freedom and resentment. In *Free Will*, edited by Gary Watson, 72–93. New York: Oxford University Press, 2003.
- Thomson, Judith Jarvis. 1975. The right to privacy. *Philosophy & Public Affairs* 4 (4): 295–314.
- Wallace, R. Jay. 1994. *Responsibility and the moral sentiments*. Cambridge, Mass: Harvard University Press.
- Watson, Gary. 1975. Free agency. *The Journal of Philosophy* 72 (8): 205–220.
- Watson, Gary. 1996. Two faces of responsibility. In *Agency and answerability: Selected essays*, 260–288. New York: Oxford University Press, 2004.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.