

The right to privacy and the deep self

LEONHARD MENGES 

University of Salzburg Faculty of Social Sciences, Austria

This paper presents an account of the right to privacy that is inspired by classic control views on this right and recent developments in moral psychology. The core idea is that the right to privacy is the right that others not make personal information about us flow unless this flow is an expression of and does not conflict with our deep self. The nature of the deep self will be spelled out in terms of stable intrinsic desires. The paper argues that this view has advantages over alternative accounts of the right to privacy, that it is extensionally adequate in interesting test cases, that there is a good reason to think that the right to privacy, thus understood, can be justified, and that this view helps identify what kind of information is protected by the right to privacy.

Keywords: the right to privacy; privacy; deep self; control; intrinsic desire; autonomy.

I. Introduction

Edward Snowden revealed that US and British secret services had access to a huge number of emails from people around the world (The Guardian 2017). The Chinese Social Credit System used surveillance drones, mobile phone apps, facial recognition systems, internet devices, and so on to evaluate the trustworthiness of Chinese citizens (Campbell 2019). The company Cambridge Analytica accessed personal information about millions of Facebook users, most of whom did not know anything about this, to develop election campaigns (The Guardian 2022). When learning about these events, a natural response is the thought that people's right to informational privacy was endangered or infringed. But was it really? And how bad would this be? To

Correspondence to: Leonhard Menges (leonhard.menges@plus.ac.at)

answer these questions, we need a clear understanding of what the right to privacy is. The goal of this paper is to provide such an understanding.

One standard account identifies the right to informational privacy with the right to some kind of control. The idea is that privacy is, at its core, about it being *up to us* what happens with personal information about us. It seems important that it is up to us what the state, certain companies, or our friends and colleagues know about us. The right to privacy, the idea goes, is there to protect this control. However, the control account has been widely criticized. Anita Allen, for example, contends ‘that the popularity of the privacy-control paradigm is problematic because there are a number of conceptual, practical, and moral limits to its plausibility’ (Allen 1999: 862; see also Thomson 1975; Parent 1983; Rickless 2007; Solove 2008; Davis 2009).

The starting point of this paper is the insight that standard control accounts have tacitly assumed that the relevant kind of control should be understood as having an effective choice about what happens with personal information. In Section II, I will argue that this and a more recent alternative view are problematic. One may respond to these problems by giving up the idea that the right to privacy is a right to control. According to alternative views, the right to privacy should be understood as the right that others cannot access or have restricted access to personal information (e.g. Allen 1988: chs 1 and 2; Davis 2009) or as the right that certain context-specific norms be respected (e.g. Nissenbaum 2009). But these views deny the seemingly intuitive idea that privacy is—at its core—about it being up to us what happens with personal information and that the right to privacy protects this control. In Section III, I will develop an account of the right to privacy that makes sense of this intuitive idea. The resulting account, that I will call the Deep Self View, is inspired by classic control views and recent developments in moral psychology. In Section IV, I will argue that the Deep Self View is extensionally adequate, normatively plausible, and explanatory powerful. Thus, I will argue that this version of the control view is, *pace* Allen, conceptually and morally plausible. Note, however, that the paper will not deal with Allen’s practicality challenge. That is, the paper will not say much about actual recent challenges to privacy. The reason for this is pragmatic. Here, I will present the core idea of an account of the right to privacy. In future work, I will test how this view works when applied to real challenges in practice.

II. Two views on the right to privacy

In this section, I will critically discuss two control accounts of the right to privacy. Most proponents of control accounts have not said much about how they understand the notion of control. Therefore, it will prove useful to constructively reinterpret the views by paying attention to other debates on

control-like concepts, especially in political philosophy, action theory, and moral psychology.

One way to flesh out the idea that we have control over something that is inspired by political philosophy is in terms of *sovereignty*. My having control over an event is my being the sovereign regarding its occurrence, it's my having 'the last word regarding the relevant matter' (Enoch 2020: 162; Enoch introduces sovereignty in his 2017). I decide whether the event occurs. I have a menu of options, choose one of them instead of the others, and my choice partly explains (in the right way) that the option is realized. For example, my control over what I do with my money consists in my being able to effectively choose between buying a new bicycle or going on vacation. My buying the bike then explains why there is a new bike in my basement.

This natural understanding of control can be plugged into the control account of the right to privacy. My right to privacy, the idea goes, is my right to effectively choose whether some other person learns about a certain piece of personal information. When strangers read my emails they, thereby, make it the case that I cannot choose anymore. As there is no option to choose from, there is nothing for me to decide. They, thereby, infringe my right to privacy. One of the authors who develop accounts along these lines is Andrei Marmor. According to him, the 'right to privacy [...] is there to protect our interest in having a reasonable measure of control over ways in which we present ourselves to others' (Marmor 2015: 13). He continues: 'your right to privacy is violated when somebody manipulates, without adequate justification, the relevant environment in ways that significantly diminish your ability to control what aspects of yourself you reveal to others' (Marmor 2015: 14; similar ideas can be found in Fried 1968: 210; Parker 1974: 286; Rachels 1975: 331; Inness 1992: 51; Rössler 2004: ch. 5.1). When the strangers read my emails, they unjustifiably and significantly diminish my ability to control how I present myself to them and others. I cannot effectively decide anymore whether they know certain information. Thereby, they violate my right to privacy, according to Marmor.

This is a natural and widely accepted understanding of the right to privacy. However, it is problematic. Consider the cases of Impeded Information Flow (see Menges 2022: 705; for similar cases, see Farber 1993; Matthews 2008; Mainz and Uhrenfeldt 2021): you write an email to your friends in a country far away telling them that you have fallen in love with *S*. This is the only way to let them know. Unbeknownst to you, your housemate hacks your email service provider (not your computer) such that every email that is sent using this service provider is encrypted in a way that neither your friends nor any other recipient can decipher it. Your housemate's only goal is to make your life complicated. Your housemate is successful: you do not succeed in telling your friends about your romantic feelings by sending them the email.

By hacking the service provider and encrypting your email, your housemate's conduct undermines your sovereignty: you cannot choose between letting your friends know about your being in love by sending them the email and not letting them know. Recall the standard control account. It says that your having a right to privacy is your having the right to effectively choose whether or not someone learns personal information about you. As your housemate unjustifiably and radically undermines your choice, the standard view implies that your housemate infringes your right to privacy. Now take Marmor again. He says that violations of your right to privacy consist in unjustified manipulations of the environment that significantly diminish your control over how you present yourself. Your housemate does exactly this. Thus, Marmor's view is committed to the claim that your housemate violates your right to privacy.

This implication is counter-intuitive. Your housemate acts problematically and diminishes your freedom and autonomy in an objectionable way. However, your housemate does not seem to infringe your right to privacy; no information is accessed, nothing is learned about you. Your housemate impedes the relevant information flow, but this does not seem to be a privacy issue. However, the standard control account implies that it is a privacy issue, which is a vice of the standard control account. Therefore, an alternative view on the right to privacy that avoids this conclusion but is, in other respects, at least equally plausible has an advantage over the standard control account.

A possible response is to revise the standard control account by specifying the thing one has control over when one has the right to privacy. One could suggest, for example, that the right to privacy is the right to be sovereign over whether personal information remains *concealed* (this view is similar but not identical to the negative control view by Mainz and Uhrenfeldt 2021; see also Véliz 2024: 214). You can choose whether a piece of information remains concealed or not and your choice partly explains what happens. The right to privacy protects this. On this view, others violate your right to privacy by unjustifiably manipulating the environment in a way that significantly diminishes your control over whether or not personal information about you remains concealed. At first sight, this may look plausible. However, this view has the same implication in cases of Impeded Information Flow that the standard control account has. Having sovereignty over concealment involves having control over *whether or not* a piece of information remains concealed. It involves that you have the last word regarding concealment. But in the case, it is not you who has the last word, it is your housemate. By concealing the information against your will your housemate significantly diminishes your control over whether or not the information that you are in love is concealed. Thus, this view also implies that your housemate infringes your right to privacy.

There is another way to spell out the right to privacy as a right to control. It does not identify control with sovereignty, but with *non-alienation* (see, again, Enoch 2017, 2020). That one has control over an event in this sense consists

in a close relation between the occurrence of the event and who one is. If the event occurs then this shows something about who one is, what one stands for, and one's commitments. I will spell out the details below, but to illustrate the idea, take my buying a new bicycle and imagine a situation in which I do so because it is very important to me to exercise. Then, buying the bike expresses something about who I am; I am not alienated from this action. Contrast this case with one in which I only buy the bike because health radicals brainwashed me. In the latter case, my buying the bike does not express my commitments and what I stand for but rather some superficial but hard-to-resist psychological mechanism. I would likely feel alienated from buying the bike when I am not in the grip of the mechanism anymore. Therefore, in an important sense I do not control buying the bicycle in the latter case, but I do control it in the case in which I buy it because I like exercising.

This understanding of control can also be plugged into the control account of the right to privacy (see Menges 2021, 2022). The result would say, roughly, that my having the right to privacy is my having the right that personal information only flows to others when this flow expresses who I am. Others infringe my right to privacy when they make the information flow even though this information flow is not based on my commitments and is, thereby, alienated from me. For example, the strangers who read my emails do so independently of what is important to me. Thereby, they infringe my right to privacy.

Some may object that non-alienation is not a kind of control and that, therefore, it is misleading to call the non-alienation account a version of the control view. Note, however, that the non-alienation account nicely makes sense of the core idea of control views that privacy is ultimately about it being up to us what happens with our personal information. The non-alienation account spells out 'up to us' in terms of commitments and who one is. But even if one rejects the idea that non-alienation is a kind of control, it is still important to find out whether the non-alienation account is plausible as an account of the right to privacy. To see if it is, note first that it has no problems with the case of Impeded Information Flow. When your housemate hacks the email service provider such that every email you send is encrypted, no information about you flows. The flow is impeded. The non-alienation account of the right to privacy says that the right to privacy is only infringed when information about you flows. Thus, the view avoids the counter-intuitive implication that your housemate infringes your right to privacy. However, the non-alienation account is underdescribed. Moreover, some suggestive ways to fill in the details are problematic. Most importantly, we need a better understanding of what it means that a person has a 'commitment' or that the person 'stands for' something.

One account fills in the details by using Frankfurtian first- and second-order desires and volitions (this idea is discussed in Menges 2021, 2022; see also Enoch 2017). On this view, I am committed to sportive cycling just in

case I have the first-order desire to do it and the second-order volition to have that first-order desire. Second-order volitions are manifestations of what Frankfurt calls ‘reflective self-evaluation’ (Frankfurt 1971: 7) and they express our rational capacities that allow us to become ‘critically aware of [our] own will’ (Frankfurt 1971: 12). The corresponding account of the right to privacy would say: my having the right to privacy is my having the right that personal information about me only flows to others if I have the first-order desire that it flows in this way and the second-order volition to have that first-order desire. Others infringe my right to privacy when they make the personal information flow even though this does not correspond to the relevant desires and volitions.

One problem for this view is that forming second-order volitions is cognitively demanding. Small children and adults with severe mental disabilities do not have many second-order volitions. Perhaps some have none because they do not have the capacity for reflective self-evaluation and cannot become critically aware of their will. The view under consideration says that the right to privacy is infringed when others let personal information flow even though this flow is not supported by the relevant pair of first- and second-order desires and volitions. It implies for persons who never have second-order volitions that their right to privacy is always infringed when others make personal information about them flow. This is because the flow is not supported by second-order volitions—they do not have second-order volitions.

This implication is counter-intuitive (see Sripada 2016 for a similar objection against Frankfurt’s account of responsibility). Imagine that my three-year-old daughter has chickenpox. I know that she loves it when her grandfather nurses her when she is sick, but she has never critically evaluated and endorsed this desire. Surely, I would not infringe her right to privacy by telling her grandfather that she has chickenpox. However, the view under consideration implies that I do. This is a problem for the view.

One may avoid this by saying that only people who have second-order volitions can have a right to privacy in the first place. This, however, is at least equally implausible. Because now one would never infringe the right to privacy of people who do not have second-order volitions—not by posting naked photos of them online, not by telling strangers intimate details about them, and so on. This is hard to accept. Thus, there is a good reason to reject the view that the right to privacy should be understood in terms of first- and second-order desires and volitions.

Many Frankfurt-inspired accounts of agency do not rely on second-order volitions. Following them, one could try to spell out non-alienation in terms of events that express our judgments about the Good (Watson 1975), our ability to transform ourselves according to what is in fact good (Wolf 1987), our concern for acting based on reasons (Velleman 1992), or our self-governing policies (Bratman 2000). However, these views are also cognitively demanding. It is far from obvious that three-year-old children and people with severe

mental disabilities have the cognitive capacities that are necessary for these agential features. But children and people with severe mental disabilities have privacy rights. Therefore, it is problematic to ground privacy rights in these cognitively demanding capacities. This is a conclusion that we should not find surprising. A primary goal of Frankfurt-inspired accounts is to make sense of free action or of what David Velleman calls ‘human action *par excellence*’ (1992: 189). However, we should not expect such an action from a three-year-old. Thus, it should not surprise us that these views do not provide us with an account of the grounds of privacy rights.

Other Frankfurt-inspired scholars have spelled out human agency in cognitively less demanding ways. I will adopt one of these ways for my positive proposal in Sections III and IV. But first, let me briefly sum up.

A standard view on the right to privacy says that it is the right to some kind of control. In this section, I have discussed two versions of this idea and I have argued that they are problematic. In contrast to other critics, however, I find the main idea very plausible that the core of privacy is that it is up to us what happens with personal information about us. In the remainder of the paper, I will present a view that captures this idea.

III. The Deep Self View on the right to privacy

The problem of the Frankfurtian version of the non-alienation account of the right to privacy is that it is cognitively demanding to have second-order volitions. However, whether some information is protected by the right to privacy does not, in most cases, depend on whether the agent has demanding cognitive capacities. Therefore, we need a less demanding account of non-alienation. Fortunately, recent literature in moral psychology offers helpful proposals. The discussion sometimes uses the notion of a person’s *deep self* to spell out what I have glossed as ‘who I am’, ‘what I stand for’, or ‘my commitments’ (the notion ‘deep self’ comes from Wolf 1987; the following presentation is mostly inspired by Arpaly and Schroeder 2014; Shoemaker 2015: ch. 1; Sripada 2016). Before presenting some details, let me put my proposal on the table. I call it the

Deep Self View: *A* has a right to privacy regarding personal information *p* just in case others have the *pro tanto* obligation directed at *A* to not make *p* flow unless (i) this flow is an expression of and (ii) does not conflict with *A*’s deep self. As long as *A* has neither waived nor forfeited the right to privacy, *B* infringes *A*’s right to privacy regarding *p* just in case *B* makes *p* flow even though it is not an expression of or conflicts with *A*’s deep self.

To illustrate, take the case again in which strangers read my emails.

First, the strangers only have a *pro tanto* obligation because it may be all-things-considered justified and, thereby, no obligation violation to read my

emails (I stipulate that right ‘violations’ are always unjustified, and that right ‘infringements’ are sometimes justified). This may be so if the strangers’ reading my emails would be the only way to prevent doomsday. Similarly, their reading my emails only infringes my right to privacy if I have *neither waived nor forfeited* it. If I have asked the strangers to read my emails, or if I am a guilty criminal and the strangers are police officers in a legitimate state, then they may not infringe my right to privacy because it has ceased to exist. In all following cases I will assume that other things are equal in the sense that the relevant agents have neither waived nor forfeited their right to privacy.

Second, the strangers’ obligation is *directed at me*. Imagine that the strangers promised their parents not to read my emails. Then, they would owe it to their parents not to read the emails. But this obligation has nothing to do with my right to privacy. My having the right to privacy says that they *owe it to me* not to read my emails.

Third, the metaphor of information flow—which is adopted from Helen Nissenbaum’s theory of privacy—invites us to think of information in analogy with liquids. The coffee in the mug on my desk does not flow anywhere. But if I take a sip, I make it flow. I cause it to move from one state into the next. Similarly, if the information that I am in love is mentioned in an unsent email on my computer, then it does not flow. If the strangers read the email and, thereby, learn about my feelings, they make the information flow. Their reading causes the information to move from one state (not being known) into the next (being known). Generally, information flow is a ‘movement, or transfer of information from one party to another’ (Nissenbaum 2004: 122) or a ‘distribution, dissemination, transmission’ of information from one party to another (Nissenbaum 2009: 145).

Fourth, what is a *deep self*? We need a plausible account of what it is to have commitments and what it is to stand for something that is not cognitively demanding. To start, note that many desires stand in an instrumental hierarchy (for an overview, see Schroeder 2020: sec. 2.2): Zeta desires to buy a new bicycle because she desires to use eco-friendly mobility. Zeta desires to use eco-friendly mobility because she desires to contribute to fighting climate change. She has this desire because she desires to leave future generations a world worth living in. Here, the hierarchy stops. Zeta does not desire to leave future generations a world worth living in because she desires something else. This is one of the things Zeta fundamentally cares about, that she desires for its own sake. Let me call a desire at the end of an instrumental hierarchy an *intrinsic desire*.

Intrinsic desires can be unstable. Imagine a situation in which people brainwash Zeta such that, for a short period, she intrinsically desires to buy a gas guzzler. But many intrinsic desires we have are stable and structure our thinking and acting in the world for a longer period. Just think of how the desire to have pleasurable experiences or that one’s children flourish shape many

people's lives. In what follows I will, unless explicitly noted, focus on stable intrinsic desires.

The claim that an agent has an intrinsic desire does not imply anything about the strength of the desire. How strong it is partly depends on how closely connected it is with different aspects of the agent's psychology (for an overview, see Schroeder 2020: sec. 2.3). There are links between intrinsic desires and the agent's motivation (Zeta is motivated to buy a new bike when her old bike breaks down), attention economy (she pays attention to news about climate change), emotions (she is worried about the lack of political action against climate change), normative judgments (she judges that fighting climate change is good), and, perhaps, other agential aspects. The stronger the desire, the more often another corresponding aspect of the agent's psychology is activated or the more intense the activation is.

Sometimes, agents' intrinsic desires fit well together when combined with certain beliefs. Imagine that Zeta also intrinsically desires animals' well-being and, therefore, desires to eat vegan food. This instrumental desire is also supported by her intrinsic desire to leave future generations a world worth living in because she believes that a vegan diet tends to be eco-friendly. Sometimes, intrinsic desires conflict. Imagine that Zeta intrinsically desires to see many different places in the world such that she instrumentally desires to fly a lot. This desire comes in conflict with her intrinsic desire to leave future generations a world worth living in because she believes that flying is counterproductive in this respect.

Intrinsic desires are not cognitively demanding. Many children and mentally disabled people stably desire things for their own sake. Intrinsic desires do not require reflection. One does not need to be aware of one's intrinsic desires—it sometimes takes years of psychotherapy to learn about one's own intrinsic desires.

Let us come back to the notion of a deep self. When we know the stable intrinsic desires of agents, then we know a lot about them. We know that Zeta has the stable intrinsic desires to leave future generations a world worth living in, support animals' well-being, and to see much of the world. These are the things she cares about, and it is an important part of who she is that she has these intrinsic desires. Thus, when an event expresses an agent's stable intrinsic desires, such as Zeta's buying a new bike, then this event is not alienated from her, but an expression of who she is. Therefore, it is natural to identify an agent's deep self with the set of her stable intrinsic desires. The more an agent's intrinsic desires mesh with other aspects of her psychology—such as other desires, motivation, emotions, attention, and normative judgments—the more robust the deep self tends to be. Nonetheless, a not-well integrated intrinsic desire that one has for a long time is still part of the agent's deep self (again, this is mainly inspired by Arpaly and Schroeder 2014; different versions of this idea can be found in Shoemaker 2015: ch. 1; Sripada 2016).

Fifth, what is it to *express* one's deep self? A natural idea is that an event only expresses an intrinsic desire when it is partly caused by it. Zeta's buying the bike needs to be etiologically rooted in her intrinsic desire to leave future generations a world worth living in to count as expressing this intrinsic desire. But not any causal relation will do. There is a causal chain from her intrinsic desire to leave future generations a world worth living in via her buying the bike to her having an accident one year later. But her having an accident plausibly does not express her intrinsic desire to leave future generations a world worth living in. There are different attempts to spell out the details of the expression relation (see, e.g. Arpaly and Schroeder 2014: ch. 3; Sripada 2016: sec. 3.2). What follows will be compatible with all of them such that we can work with an intuitive understanding of 'event *E* expresses desire *D*'. More importantly for this context, sometimes the actions of other people can express one's deep self. Imagine that my desire that my friend learn about my being in love with *S* is rooted in an intrinsic desire and does not conflict with another one. I, therefore, send my friend an email telling her about my feelings. This causal chain results in my friend's learning about my being in love. That my friend learns about my feelings is an action of another agent that expresses and does not conflict with my deep self.

Sixth, an event *conflicts* with an agent's deep self just in case it frustrates a stable intrinsic desire of the agent. When I stably and intrinsically desire that others not pity me and, therefore, instrumentally desire that they not know about my hard childhood, then someone's telling them about my hard childhood conflicts with my deep self. It is possible for an event to express an agent's deep self *and* conflict with it. When Zeta flies to New Zealand for vacation, then this expresses her desire to see as much of the world as possible. But it is also in conflict with her desire to not contribute to climate change. Thus, this event expresses and conflicts with Zeta's deep self (I will discuss an analogous privacy case in a moment).

Now, we have a new account of the right to privacy on the table: it says that our right to privacy regarding a certain piece of personal information is our right that this piece of information only flows when the flow expresses and does not conflict with our deep self. Thus, the right to privacy essentially protects our deep self as the source of flows of personal information about ourselves.

Let me discuss three critical responses. The first says that the Deep Self View is unnecessarily complex because clause (i) is not needed. On the alternative view, my right to privacy is my right that others do not make personal information about me flow unless it does not conflict with my deep self. This view does not require that the flow be an expression of my deep self.

This view is more elegant than the Deep Self View. However, it is too permissive. To illustrate, imagine that none of my intrinsic desires conflicts with my colleagues' learning about my hard childhood. Intuitively, they would

still infringe my right to privacy when they put their ear to the keyhole and, thereby, learn about my childhood when I talk to my partner about it behind closed doors. The simpler view comes in conflict with this intuition. The Deep Self View has the intuitively plausible implication. It says that my colleagues would *not* infringe my right if their putting their ears to the keyhole is caused in the right way by one of my intrinsic desires—perhaps, they know that I deeply want them to learn about my childhood and that I'm only talking to my partner to let them know about it.

The second response says that the Deep Self View is not informative. Let me elaborate. When you own a computer, then you can waive your property right by giving it to your neighbor. When you have a right to privacy regarding health information, then you can waive it by telling your doctor. Some may say that there is not a big difference between waiving a right and expressing one's deep self in certain ways. The idea is that expressing your deep self by telling the doctor about your health problems and waiving your right to privacy regarding the information are extremely similar or even identical. Then, the Deep Self View would not be interesting because it would say that your right to privacy regarding the flow of a certain piece of personal information is others having the pro tanto obligation directed at you to not make it flow unless you waive that right. However, this is not informative. Obviously, people can waive their right to privacy. We wanted to know what the right to privacy is, independent of its being possible to waive it.

The objection rests on the idea that waiving rights and expressing one's deep self in certain ways are very similar. However, this is not so. One difference is that the deep self consists in one's *stable* intrinsic desires. These are intrinsic desires that structure one's thinking and acting for a longer period. However, waiving one's right does not need to express a stable intrinsic desire. If I want to be kissed by *S*, if I permit *S* to kiss me, if I know what is going on, and if I'm not drunk, manipulated, and so on, then I, thereby, waive my right not to be kissed. This is so even if my desire is not stable, and I developed it just before and stopped having it right after the kiss. Thus, I can waive rights without, thereby, expressing my deep self.

It is important for us to be able to waive rights even though this would conflict with our deep self. Imagine that Yella stably and intrinsically wants to fight discrimination against homosexuals, which is why she instrumentally desires to make her own homosexuality public and present it as the normal thing she wants it to be. But Yella also stably and intrinsically desires to not be the victim of homophobic discrimination, which is why she instrumentally desires to not make her homosexuality public. If Yella's friend told others about her homosexuality to help fulfill her desire to fight discrimination, this would express one part of her deep self and conflict with another part. Intuitively and according to the Deep Self View, this would infringe her right to privacy. In situations of this kind, it is important for us to be able to waive our rights

even if this comes in conflict with some of our stable intrinsic desires. If Yella waives her right to privacy, then others will not infringe it when they let information about her homosexuality flow even though it conflicts with part of her deep self.

This brief discussion of waiving rights may suggest a third response to the Deep Self View (thanks to an anonymous referee for the following suggestion). Often, waiving rights does not need a demanding kind of agency. Perhaps, having an unmanipulated, uncoerced wish and a certain understanding of the situation are sufficient for being able to waive the right not to be kissed by *S*. Let us call the relevant agency, regardless of what, exactly, it consists in, Weak Agency. Why not, the response continues, spell out the right to privacy in terms of Weak Agency instead of intrinsic desires? The resulting Weak Agency View would say that I have a right to privacy regarding personal information just in case others have the pro tanto obligation directed at me to not make it flow unless (i) this flow is an expression of and (ii) does not conflict with my Weak Agency.

It seems that the Weak Agency View is too permissive. Consider a case that is a bit more fantastic than those discussed so far: my best friend knows that I, generally, want to keep it secret that I am in love with *S* and that this is based on some of my stable intrinsic desires. Imagine further that because of some fantastic event—perhaps a lightning strike—I lose the relevant intrinsic desires for a short period and, instead, wish that everyone learns that I am in love with *S*. Imagine that I fulfill the requirements for Weak Agency in that period, that I do not explicitly waive any rights, and that, after the period, I will have the intrinsic desires again that make me want to keep the information secret. Finally, imagine that my friend knows all this. The Weak Agency View implies that my friend would not infringe my right to privacy, when, in order to fulfill my wish, she tells everyone that I am in love with *S*.

I find this counter-intuitive. If I do not waive my right and if my wish that everyone learns about my feelings is not based on stable intrinsic desires, it seems more plausible to me that my friend would infringe my right to privacy. More generally, given the complexity of our social lives, we need quick and easy ways to waive certain rights, for example when we let someone go first in a queue or leave our seat to another passenger. For this to work, a superficial kind of agency seems sufficient. It seems intuitive to me, however, that the right to privacy protects the fact that it is up to us what happens with personal information about us in a deeper sense. The Deep Self View spells out the details of this intuitive idea, the Weak Agency View does not. Obviously, this is not a knock-down argument. Given the overall goals of this paper, however, I must leave the Weak Agency View for another occasion.

To sum up, the Deep Self View says that the right to privacy protects our deep self as the source of flows of personal information about us. In what follows, I will show that there are good reasons to accept this view. I will do so

by arguing that the Deep Self View is extensionally adequate in some interesting cases (Section IV.1), normatively plausible (Sections IV.2 and IV.3) and explanatorily powerful (Section IV.4).

IV. Reasons to accept the Deep Self View

IV.1 Extensional adequacy

In Section II, I suggested that standard control views and the Frankfurt-inspired non-alienation account are extensionally inadequate. In other words, they identify the wrong actions as infringements and as non-infringements of the right to privacy. I will now test whether the Deep Self View meets the challenge of extensional adequacy by applying it to some potentially problematic cases.

First, recall the cases we have already come across. The strangers who read my emails without having asked me circumvent my intrinsic desires. My deep self plays no role in their making the information flow. Therefore, the view implies that they infringe my right to privacy. Now, take the case of Impeded Information Flow where my housemate hacks the email service provider and encrypts my email without reading it. As my housemate does not make the information flow—the flow is impeded—my right to privacy is not affected. Finally, recall my three-year-old daughter who has chickenpox and wants her grandfather to take care of her but has not critically endorsed this desire. Plausibly, my telling her grandfather that she has chickenpox is partly causally rooted in and, thereby, expresses her intrinsic desires, such that I do not infringe her right to privacy. These are the intuitively correct results, which is a good starting point for the Deep Self View.

Now consider a case that is often presented as a problem for control views, the Threatened Loss Case: ‘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’ (Thomson 1975: n. 1). The assumption is that, intuitively, the neighbor does not infringe my right to privacy if the device is not turned on, used, and personal information about me is not accessed. The Deep Self View has this implication. It says that the stranger would infringe my right to privacy by making the information flow when this is not an expression of my stable intrinsic desires. However, the stranger does not make the information flow if the device is not turned on. The neighbor does not learn anything about me. Thus, the Deep Self View has the intuitively correct implication.

More generally, the Deep Self View has the intuitively correct implications in important test cases.

IV.2 Privacy and autonomy

A classic worry about control accounts of the right to privacy is that the right, thus understood, cannot be justified. On this view, it is simply false that we (typically) have the right to control information about us (Allen 1999: sec. 4). In this and the following subsection, I will argue that there are good reasons to believe that we have a right to privacy as it is identified by the Deep Self View.

One way to justify the right to privacy, as it is presented in this paper, is to ground it in the value of personal autonomy (for another attempt to ground the right to privacy in the value of autonomy, see Rössler 2004: ch. 3; 2017; for critical discussion, see Mokrosinska 2018). By ‘personal autonomy’ I do not mean the Kantian idea of being subject to laws that one has legislated oneself or the kind of agency that is relevant for free will. Rather, I mean the idea that people are the authors of their own life stories and that it is a good thing to be the author of one’s own life story (Raz 1988: 369). The idea is, roughly, that an event that concerns oneself is determined by oneself and what one cares about. I am autonomous regarding an event insofar its occurrence is a matter of what I stand for and not completely determined by something alien to me (see Enoch 2020: 162). My buying the bike because I like exercising is autonomous, whereas my buying it because I was brainwashed by health radicals is not.

The question of how to spell out this kind of autonomy can be answered in many ways (for an overview, see Christman 2020). But one way should be obvious from what has been said above, namely in terms of the agent’s deep self. According to this view, if an event expresses an agent’s deep self, then the agent is autonomous regarding this event. That is, when my buying the bike expresses my stable intrinsic desires and, thereby, my deep self, then my buying the bike is an exercise of my autonomous agency. In an intuitive sense, it depends on what is important to me such that the act is another line in the story of my life, written by me. Most adults, many children, and many persons with cognitive disabilities can be autonomous and authors of their life stories in this cognitively undemanding sense (see, e.g. Mullin 2014; Hannan 2018; Matthews and Mullin 2023).

Perhaps, different kinds of agents differ in how robust their deep selves and, thereby, their autonomous agency typically are. It seems plausible, that the intrinsic desires of young children, for example, are less well integrated with other aspects of their psychology than the intrinsic desires of some adults.¹ Recall, however, that a less-well integrated and even isolated intrinsic desire

¹This may be used to explain why certain rights sometimes differ in their strength. One could suggest that a more robust autonomous agency is normatively more important such that it grounds stronger rights than a less robust autonomous agency. In what follows, I will leave the strength of rights aside because an adequate discussion would lead me too far away from the

that one has for a long time is still part of one's deep self. If an event manifests such a desire, then the agent is still autonomous regarding the event because it is still *her* intrinsic desire that determines it.

Based on this proposal, one can draw a direct line from autonomy to the right to privacy. The Deep Self View says that the right to privacy is the right that flows of personal information be grounded in the agents' deep self. If others make the information flow even though this is not an expression of or conflicts with the agents' deep self then, other things being equal, the agents' right to privacy is infringed. Combine this with the proposal I have just made, that if an event expresses the agents' deep self, then the agents are autonomous regarding this event. Then, the right to privacy also protects personal autonomy. By protecting the agents' deep self as the source of information flows, it protects the agents' being autonomous regarding information flows. To illustrate, when the strangers read my emails, they make it the case that the information flow is not an expression of my deep self and, thereby, they interfere with my personal autonomy. My right to privacy protects me from this interference. And if personal autonomy is so valuable that it grounds rights that protect it, then we have a right to privacy. But is this cognitively undemanding kind of autonomy really so valuable? Answering this question in detail would go beyond the scope of the paper. But let me, nonetheless, try to make an affirmative answer plausible.

It seems plausible (at least to many) that what a person stands for and her commitments deserve, other things being equal, respect. Imagine that you believe that your friend would be a perfect judge and a mediocre philosopher. But she stably and intrinsically desires to be a philosopher and does not desire to be a judge. Other things being equal, you should not interfere with the decisions that express her commitments. You may discuss the issue with her, but you surely should not secretly delete her application letters for graduate programs in philosophy. Such an interference would be disrespectful. This is so even if you are right and she would be the perfect judge but a mediocre philosopher. Similarly, when your friend stably and intrinsically does not want others to know about her hard childhood, then, other things being equal, you should not tell others about it. Again, you may discuss it with her. But it would be disrespectful to disregard her intrinsic desires even if it would, in fact, be best for her if more people knew about her childhood. In both cases the value of living out one's intrinsic desires corresponds to others' pro tanto obligation not to interfere. And this can be spelled out in terms of the value of autonomy. As David Enoch puts it:

Respecting people's desires and commitments is sometimes a way of responding appropriately to their autonomy—to the fact that, at least often, being the author of one's

goals of the paper. I thank an anonymous referee for pressing me on discussing the robustness of the deep self.

own life story is an aspect of living a good life, and perhaps also that people have a right to be—to an extent—the authors of their life story (Enoch 2017: 25).

What seems plausible for desires in general is particularly plausible for an agent's stable intrinsic desires. The right to privacy, as it is presented by the Deep Self View, protects in an almost literal sense the agents' ability to write their own life stories. It makes sure that what others 'read' about them depends on their stable intrinsic desires.

Some may object that a valuable kind of autonomy is not as closely connected with the agent's deep self as I have suggested. Perhaps, they would say that valuable autonomy is cognitively more demanding than having intrinsic desires. Therefore, I will now present three considerations in favor of the claim that it is in our fundamental interest that personal information about us only flow if this flow expresses and does not conflict with our deep self.

IV.3 The instrumental, representative, and symbolic values protected by privacy

First, it is often *instrumentally* valuable for me that what happens to me depends on what I intrinsically desire (see Scanlon 1998: 252 for a similar idea; see also Munch 2021: 3786–87). If I intrinsically desire to experience the taste of fish, then it is often helpful that what I get on my plate depends on this desire. For otherwise I would often end up eating food that I do not like, which would, other things being equal, be bad for me. Similarly, if I stably and intrinsically desire that people do not pity me and, therefore, instrumentally desire that others not learn about my hard childhood, then it is often helpful when the relevant information flows depend on my desires. That others have the pro tanto obligation to not make information about my hard childhood flow makes it more likely that others will not pity me, which is, given my desires and other things being equal, good for me. Thus, the right to privacy as it is presented by the Deep Self View is supported by instrumental considerations. Other things being equal, the right serves my fundamental interest that good things happen to me and that bad things do not happen to me.

Second, sometimes it is of *representative* value that an event expresses my deep self (see Scanlon 1998: 252–53 for a similar idea). Imagine that I need to buy a gift for my son. I can do so by combining the top-ten list of gifts for thirteen-year-olds and a random generator that picks one option from the list. But I can also do so by asking myself what I intrinsically desire for him and then see which gift may help him get it. Perhaps I would buy the same thing, but the gift would have another meaning and, thereby, another representative value. The special value of the gift is explained by the fact that it expresses what I intrinsically desire for my son. Similarly, imagine that Xenia intrinsically and stably desires to have intimate bonds with her family

and, therefore, instrumentally desires to personally tell her close relatives that she is pregnant. She starts with her brother. Consider the following two cases: in one, the brother calls their parents, without regard for Xenia's desires, and tells them that she is pregnant before Xenia can. In the other, it is Xenia who first tells them that she is pregnant. Furthermore, imagine that in both cases the same pieces of information about the pregnancy flow to the parents. There is still an important difference between the two cases. This is because it sometimes matters who the source of an information flow is. In the second case, the information flow is an expression of Xenia's stable and intrinsic desire to have intimate bonds with her family, shaped by herself. Therefore, the information flow has a special meaning and representative value.

The same is true for children. Parents of children close in age may be familiar with situations like the following: something important happened to child *A*, who wants to tell her parents about it, but child *B* is faster and passes on the information. This is a common cause of fights. It was important for *A* to be the source of the information flow partly because it helps to shape her relationship with her parents. The idea generalizes: it is often in our fundamental interest that flows of personal information express our deep self because such a flow can have a specific representative value. The right to privacy, as it is presented by the Deep Self View, protects this interest (the idea that privacy enables us to shape relationships can be found in, e.g. Fried 1968; Rachels 1975; Gerstein 1978; Moore 2010: ch. 2; Marmor 2015).

Third, it is sometimes of *symbolic* value that what happens with personal information about me is a matter of my stable intrinsic desires (see Scanlon 1998: 253 for a similar idea). To see this, note that intrinsic desires do not always ground obligations or reasons of the same weight. When my three-year-old daughter stably, intrinsically, and strongly desires that her teeth not be brushed, then there is a reason of some weight not to brush her teeth. But it seems that prudential considerations easily override this reason. Now imagine that my twenty-year-old daughter has the same desire of the same strength. Then I, surely, have a much stronger obligation not to brush her teeth. A plausible explanation for this is that the stable intrinsic desires of competent adults are, other things being equal, normatively more relevant than the stable intrinsic desires of a three-year-old. If I would brush the teeth of my adult daughter against her will, then this would be—among other things—demeaning. It would express that I do not regard her as a competent adult. The same is true for information flows. Perhaps, my three-year-old daughter really wants me not to talk to the doctor about her chickenpox, which is grounded in some of her stable intrinsic desires. This gives me some reason not to do it. But it is easy to imagine situations in which the reason is overridden by considerations about the value of her long-term health such that I can permissibly infringe her privacy right. But if I talk to the doctor about the health problems of my twenty-year-old daughter even though she stably and intrinsically wants

me not to do it, then this would show that I do not regard her as a competent adult. More generally, flows of pieces of personal information about us sometimes have a symbolic value. They express how we are regarded in our community. Therefore, it is sometimes important that it depend on our stable intrinsic desires whether a piece of personal information flows. It protects our standing as competent members of our community. The right to privacy as it is presented by the Deep Self View protects this value.

To sum up, there are good reasons to believe that it is often in our fundamental interest that personal information about us only flow if this expresses and does not conflict with our deep self. Most generally, this is because it is then, to a certain degree, up to us how we relate to others. On the assumption that our rights protect our fundamental interests, one should conclude that we often have the right to privacy as the Deep Self View presents it. Moreover, there is a good reason to think that the right to privacy, thus understood, protects an important kind of personal autonomy. Taken together, the prospect for justifying the right to privacy, as it is presented by the Deep Self View, is quite good.

IV.4 Personal information and ignorance

The ideas presented in the preceding subsection help answer two important questions: first, what information is protected by the right to privacy? Second, what should we do when we don't know whether another person has the stable intrinsic desire to let a piece of personal information about her flow?

So far, the Deep Self View says that *personal information* is protected by the right to privacy. This must be information about a person. That I have two hands is information about me. However, you surely would not infringe my right to privacy by letting this piece of information flow without this being an expression of my deep self by, for example, telling your friend that I have two hands. This raises the question of what information is such that it can be protected by a person's right to privacy.

The preceding subsection helps answer this question. The idea is to look at when it is instrumentally, representatively, or symbolically valuable for agents that their deep selves be the sources of flows of information about them. The information that I have two hands is such that it is in none of these ways valuable for me that my deep self is the source of its flow. In my current context, no bad things will happen to me when others learn that I have two hands, I can't use this information to shape personal relationships, and it does not express anything about my standing in my community when people talk about my having two hands. Therefore, the suggestion goes, this piece of information is not *personal* in the sense at issue: it is not protected by my right to privacy. More generally, it must be in the context at issue instrumentally, representatively, or

symbolically good for me that my deep self is the source of an information flow for its being true that the information is protected by my right to privacy.

An interesting upshot of this view is that whether a piece of information is protected by the right to privacy is in a plausible way context-dependent. Imagine that we do not know anything about the context of some people and ask what information about them is protected by the right to privacy. Some may say that the right protects *every* piece of information about them, and others may say that it protects *no* information at all. My alternative proposal is that, without specifying the context, there is no answer to the question of whether a certain piece of information is protected by the right to privacy. We need to know whether it is in the relevant context instrumentally, representatively, or symbolically valuable that a piece of information only flows if this flow expresses the agent's deep self. The reason for this is that, according to the ideas presented in Section IV.3, the right to privacy fundamentally protects our control over how we relate to others. And how we can relate to others depends on the social context we live in.

To illustrate, take the information that someone has two hands. Without any context, I suggest, there would be no answer to the question of whether it is protected by the right to privacy. In my current context, this piece of information about me is not protected because it is irrelevant for me whether my deep self is the source of the flow of this piece of information. But imagine a context in which agents with two hands are forced to fight in an unjust war. Then, it is of high instrumental importance that the information that one has two hands only flows if this is an expression of one's deep self. Therefore, this piece of information is protected by the right to privacy in this context.²

The second question left open starts from the observation that it is sometimes hard to know what other people's stable intrinsic desires are. In such cases, we do not know whether a flow of personal information would express or conflict with the person's deep self. What should we do then?

The preceding considerations support the following suggestion: if the piece of information about a person is such that, given the relevant context, it is instrumentally, representatively, or symbolically valuable for her that her deep self is the source of the information flow, then we should, other things being equal, not let it flow unless we have a sufficient reason to believe that the flow would express and not conflict with her deep self. Imagine that you

²Note that the Deep Self View also explains why my right to privacy sometimes protects information that I am not aware of. Imagine that my heritage is such that if certain people learn about it, they will discriminate against me. However, I don't know anything about my heritage. If another person learns about my heritage and tells others about it, then this person would make a piece of information flow that is such that it is valuable for me that it only flows when this expresses my deep self. And in the case at issue, it does not express my deep self. Thereby, the person infringes my right to privacy, according to the Deep Self View. Thanks to an anonymous referee for pressing me on this issue.

know about a friend's homosexuality, but you don't have good evidence for her stable intrinsic desires regarding who learns about it. If you live in a homophobic society, then the context is such that it is of instrumental value for your friend that her deep self is the source of the relevant information flow. Then you should not let the piece of information flow because doing so would risk violating your friend's right to privacy.

Let me briefly sum up. First, the Deep Self View is extensionally adequate in interesting test cases; second, the chances of justifying the right identified by the Deep Self View look good; third, the view has resources to identify many kinds of information that are not protected by the right to privacy; and fourth, it suggests what to do when we are uncertain about another person's deep self. Considered together, I take these to be good reasons to accept the Deep Self View.

V. Conclusion

This paper has presented the Deep Self View on the right to privacy. According to it, the right to privacy is the right that others do not make personal information about us flow unless this flow is (i) an expression of and (ii) does not conflict with our deep self. Now, what does this view say about the Snowden revelations, the Chinese surveillance system, and Cambridge Analytica? The details need to be spelled out elsewhere, but the core idea would be this: if these institutions make personal information about an agent flow regardless of the agent's deep self, then they infringe the agent's right to privacy. And if they risk making the information flow, then they endanger the agent's right to privacy. As the right to privacy protects fundamental interests, both would be very bad.

Acknowledgment

I am very grateful for written comments by Hannah Altehenger, Lauritz Munch, and Lizzy Ventham and for helpful conversations with Munch, Jakob Mainz, and Björn Lundgren in Aarhus (June 2022) and with Asja Ahatović, Leonie Eichhorn, Lizzy Ventham, and Maria Seim in Salzburg (November 2023). I presented this paper in Salzburg (November 2021), Frankfurt a. M. (May 2022), Aarhus (June 2022), Berlin (conference of the Society for Analytic Philosophy, September 2022), Münster (December 2022), and Munich (June 2024), and I thank the audiences for constructive comments and questions. I am very grateful to the anonymous referees who helped to make the paper better. Thanks also to Claire Davis (and Lizzy Ventham) for proofreading. This research was funded in whole or in part by the Austrian Science Fund (FWF) 10.55776/P 36226-G. For open access purposes, the author has applied

a CC BY public copyright license to any author accepted manuscript version arising from this submission. Work on this paper is part of the research project “The Source View on the Right to Privacy”.

References

- Allen, A. L. (1988) *Uneasy Access: Privacy for Women in a Free Society*. Totawa, NJ: Rowman & Littlefield.
- Allen, A. L. (1999) ‘Privacy-as-Data Control: Conceptual, Practical, and Moral Limits of the Paradigm’, *Connecticut Law Review*, 32: 861–75.
- Arpaly, N. and Schroeder, T. (2014) *In Praise of Desire*. New York: OUP.
- Bratman, M. E. (2000) ‘Reflection, Planning, and Temporally Extended Agency’, *The Philosophical Review*, 109/1: 35–61. <https://doi.org/10.2307/2693554>
- Campbell, C.h. (2019) ‘How China Is Using Big Data to Create a Social Credit Score’, *Time*. <https://www.time.com/collection/davos-2019/5502592/china-social-credit-score/>, accessed 24 July 2024.
- Christman, J. (2020) ‘Autonomy in Moral and Political Philosophy’, in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*. Fall 2020. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2020/entries/autonomy-moral/>, accessed 24 July 2024.
- Davis, S. (2009) ‘Is There a Right to Privacy?’, *Pacific Philosophical Quarterly*, 90/4: 450–75.
- Enoch, D. (2017) ‘Hypothetical Consent and the Value(s) of Autonomy’, *Ethics*, 128/1: 6–36. <https://doi.org/10.1086/692939>
- Enoch, D. (2020) ‘False Consciousness for Liberals, Part I: Consent, Autonomy, and Adaptive Preferences’, *The Philosophical Review*, 129/2: 159–210. <https://doi.org/10.1215/00318108-8012836>
- Farber, D. A. (1993) ‘Book Review: Privacy, Intimacy, and Isolation by Julie C. Inness’, *Constitutional Commentary*, 510/10: 510–9.
- Frankfurt, H. (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>
- Fried, C.h. (1968) ‘Privacy [A Moral Analysis]’, in F. D. Schoeman (ed.) *Philosophical Dimensions of Privacy: An Anthology*, pp. 203–23. Cambridge: CUP.
- Gerstein, R. S. (1978) ‘Intimacy and Privacy’, *Ethics*, 89/1: 76–81.
- Hannan, S. (2018) ‘Childhood and Autonomy’, in G. Calder, A. Gheaus and J. de Wispelaere (eds) *The Routledge Handbook of the Philosophy of Childhood and Children*, pp. 112–22. London: Routledge.
- Inness, J. (1992) *Privacy, Intimacy, and Isolation*. New York: OUP.
- Mainz, J. T. and Uhrenfeldt, R. (2021) ‘Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy’, *Res Publica*, 27/2: 287–302. <https://doi.org/10.1007/s11158-020-09473-1>
- Marmor, A. (2015) ‘What Is the Right to Privacy?’ *Philosophy & Public Affairs*, 43/1: 3–26.
- Mathews, G. and Mullin, A. (2023) ‘The Philosophy of Childhood’, in E. N. Zalta and U. Nodelman (eds) *The Stanford Encyclopedia of Philosophy*. Fall 2023. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2023/entries/childhood/>, accessed 24 July 2024.
- Mathews, S. (2008) ‘Privacy, Separation, and Control’, *The Monist*, 91/1: 130–50. <https://doi.org/10.5840/monist200891116>
- Menges, L. (2021) ‘A Defense of Privacy as Control’, *The Journal of Ethics*, 25/3: 385–402. <https://doi.org/10.1007/s10892-020-09351-1>
- Menges, L. (2022) ‘Three Control Views on Privacy’, *Social Theory and Practice*, 48/4: 691–711. <https://doi.org/10.5840/soctheorpract2022811172>
- Mokrosinska, D. (2018) ‘Privacy and Autonomy: On Some Misconceptions Concerning the Political Dimensions of Privacy’, *Law and Philosophy*, 37/2: 117–43. <https://doi.org/10.1007/s10982-017-9307-3>
- Moore, A. (2010) *Privacy Rights: Moral and Legal Foundations*. University Park, PA: The Pennsylvania State University Press.

- Mullin, A. (2014) 'Children, Paternalism and the Development of Autonomy', *Ethical Theory and Moral Practice*, 17/3: 413–26. <https://doi.org/10.1007/s10677-013-9453-0>
- Munch, L. (2021), 'Privacy Rights and 'Naked' Statistical Evidence', *Philosophical Studies*, 178/11: 3777–95. <https://doi.org/10.1007/s11098-021-01625-0>
- Nissenbaum, H. (2004) 'Privacy as Contextual Integrity', *Washington Law Review*, 79/1: 119–58.
- Nissenbaum, H. (2009) *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford, CA: SUP.
- Parent, W. A. (1983) 'Privacy, Morality, and the Law', *Philosophy and Public Affairs*, 12/4: 269–88.
- Parker, R. B. (1974) 'A Definition of Privacy', *Rutgers Law Review*, 27/2: 275–97.
- Rachels, J. (1975) 'Why Privacy Is Important', *Philosophy & Public Affairs*, 4/4: 323–33.
- Raz, J. (1988) *The Morality of Freedom*. New York: OUP.
- Rickless, S. C. (2007) 'The Right to Privacy Unveiled', *San Diego Law Review*, 44/1: 773–99.
- Rössler, B. (2004) *The Value of Privacy*. Cambridge: Polity Press.
- Rössler, B. (2017) 'Privacy as a Human Right', *Proceedings of the Aristotelian Society*, 117/2: 187–206. <https://doi.org/10.1093/arisc/aox008>
- Scanlon, T. M. (1998) *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Schroeder, T. (2020) 'Desire', in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*. Summer 2020. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/sum2020/entries/desire/>, accessed 24 July 2024.
- Shoemaker, D. (2015) *Responsibility from the Margins*. New York: OUP.
- Solove, D. J. (2008) *Understanding Privacy*. Harvard University Press.
- Sripada, C.h. (2016) 'Free Will and the Construction of Options', *Philosophical Studies*, 173/11: 2913–33. <https://doi.org/10.1007/s11098-016-0643-1>
- The Guardian (2017) 'The NSA Files', <http://www.theguardian.com/us-news/the-nsa-files>, accessed 1 July 2024.
- The Guardian (2022) 'The Cambridge Analytica Files', <https://www.theguardian.com/news/series/cambridge-analytica-files>, accessed 1 July 2024.
- Thomson, J. J. (1975) 'The Right to Privacy', *Philosophy & Public Affairs*, 4/4: 295–314.
- Véliz, C. (2024) *The Ethics of Privacy and Surveillance*. New York: OUP.
- Velleman, J. D. (1992) 'What Happens When Someone Acts?' in J. M. Fischer and M. Ravizza (eds) *Perspectives on Moral Responsibility*, pp. 188–210. Cornell University Press. <https://doi.org/10.7591/9781501721564-008>
- Watson, G. (1975) 'Free Agency', *The Journal of Philosophy*, 72/8: 205–20.
- Wolf, S. (1987) 'Sanity and the Metaphysics of Responsibility', in G. Watson (ed.) *Free Will*, pp. 372–87. New York: OUP.