

Why the NSA didn't diminish your privacy but might have violated your right to privacy

Lauritz Aastrup Munch, Aarhus University, laumu@cas.au.dk

Penultimate version (pre-proof, final published version will differ slightly). Final version forthcoming in *Analysis*. Please cite the published version when possible.

1. Introduction

Some believe that privacy is the absence of epistemic access to some fact about you (e.g., Matheson 2007; Blaauw 2013; Fallis 2013; Munch and Mainz 2023; Véliz 2024). To illustrate, why is it plausible to think that you have your privacy diminished when someone reads your diary? The epistemic account suggests that it is because of the reader learning facts about you from reading it.

The epistemic account is plausible but has a problem: it aligns poorly with judgments about what types of acts that can violate the *right to privacy*. The challenge is well illustrated in Macnish's (2018) influential discussion of mass surveillance, prompted by Edward Snowden in 2013 revealing the existence of the NSA's so-called PRISM program. The PRISM program involved large-scale and indiscriminate collection of internet and mobile phone data (Menges 2020). Intuitively, it seems right to say that the NSA did violate, or at least could have violated, the privacy rights of people who had their information collected via the PRISM program (Macnish 2018; Lundgren 2021; Véliz 2024). And yet, if we take seriously the epistemic account of privacy, it's not clear how that could follow. As Macnish points out, for the most part, the NSA only *collected* information about people. For the most part, they didn't access – i.e., have a look at - this information. But if privacy is diminished only if a certain sort of epistemic access occurs, then it would seem like the NSA didn't, for the most part, diminish the privacy of the people who had their information collected.

Macnish concludes that the NSA didn't, for the most part, violate the privacy rights of people who had their information collected. Perhaps the NSA acted wrongly, but the wrongdoing was of a different kind. However, some have taken issue with this conclusion because it seems intuitive that the NSA did violate privacy rights *and* did so by collecting information (e.g., Lundgren 2021; Véliz 2024: 147). This creates a challenge for the epistemic account of privacy: explain how the right to privacy can be violated by information collection and do so in a way that doesn't put pressure on the idea that privacy, fundamentally, is a function of epistemic access. More directly:

explain how to violate someone's privacy right without affecting their privacy. Call this *the alignment challenge*.

The purpose of this paper is to answer the alignment challenge. I start by criticizing two recent attempts at answering it, one focusing on risk (Lundgren 2021), and another focusing on modal robustness (Véliz 2024). Next, I develop my preferred response in the form of an account according to which the set of activities that may violate the right to privacy can be given a unified description as instances of a general type of activity known as 'inquiry'. If acts that violate someone's right to privacy characteristically involve inquiry, is no surprise that the right to privacy and privacy could come apart, since one can initiate inquiry without completing it, and initiate inquiry that results in false beliefs.

2. The Alignment Challenge

Intuitively, you lose privacy when: a peeping tom observes you in the shower, someone reads your diary or a friend shares your secret (e.g., Thomson 1975; Marmor 2015). What do these cases have in common? Proponents of so-called *access* accounts of privacy say that you lose privacy when, and because, someone else accesses some object of yours (Véliz 2024). Proponents of the *epistemic* access account of privacy say that this object consists of truthful propositions about you (facts) and 'access' means *knowing* these propositions:

The Epistemic Account of Privacy (EAP): A has privacy with regards to a fact about them and with regards to B just when B doesn't know (or rationally believe) this fact.

EAP delivers attractive verdicts in a range of cases. Why do you lose privacy when somebody observes you in the shower? *Answer:* someone comes to know what you look like naked. Why do you lose privacy when your secret is passed on? *Answer:* someone now knows what you meant to conceal. There is some disagreement over the details of the relevant epistemic relation, but we can set this aside here (Fallis (2013), Blaauw (2013), Matheson (2007), Le Morvan (2015; 2018), Véliz (2024).

It's natural to suspect that there would be a close connection between privacy and the right to privacy to the effect that violations of the right to privacy must, in some sense, affect privacy (Lundgren 2020; Menges 2021). This suggestion poses a challenge to EAP since it seems to allow for the possibility of cases that (intuitively) involve violations of the right to privacy but where privacy remains unaffected. Consider two such:

Incomplete Investigation. Barrack covertly sets up a camera in Dough’s house. The camera collects information about Dough, but Barrack never looks at the footage.

Mistaken Conclusion. Shoshana reads in Linda’s diary. However, Shoshana is a bad reasoner and forms only false beliefs about Linda despite the evidence being of high quality.

In these cases, intuitively, Dough’s and Linda’s right to privacy is being violated. However, EAP does a bad job of making sense of this. In *Incomplete Investigation* Barrack never looks at the footage, so the access account implies that Dough’s privacy remains unaffected. Moreover, because privacy is normally taken to be factive, Linda’s privacy remains unaffected in *Mistaken Conclusion* because Shoshana forms false beliefs only.¹ Since the challenge is that of explaining how a plausible account of what privacy is (EAP) can be used to deliver a plausible account of when the right to privacy is being violated, we can refer to the challenge as the *alignment challenge*.

Access theorists have tried a couple of different responses. Lundgren (2021) says that the right to privacy is broader than the concept of privacy because the best account of rights says that you may violate someone’s right to something (e.g., privacy) if you act in a way that *risks* affecting this something (e.g., privacy). Risk, possibly, allows us to make sense of the cases above without casting doubt on EAP.

There are at least two problems with Lundgren’s proposal. First, it requires us to take a controversial major premise about the scope of rights on board. As he suggests, “the modification (...) recognizes that the right to privacy includes a right not to have others put one’s privacy at substantial risk.” (Lundgren 2021: 385-386). This premise may be questioned, but even if true, we might prefer a solution more localized to the topic at hand. Second, the solution is insufficiently general. We can imagine cases where the risk of losing privacy is negligible but where this doesn’t affect the intuition that the right to privacy has been violated. Suppose, for instance, in *Mistaken Conclusion*, that the risk of Shoshana learning facts about Linda is trivial or close to zero (she’s that bad of a reasoner). Even still, it seems that Shoshana violates Linda’s right to privacy.

In her recent book, Carissa Véliz offers a different take. She says that what privacy is and what the right to privacy protects against come apart because the right to privacy is a “right to a rich or robustly demanding good” (Véliz 2024: 145; compare Pettit 1997).² Véliz says that the right to privacy protects privacy (as characterized by EAP), but that it reaches beyond this. As she puts it, “What we informally call “the right to privacy” [...] is more precisely a right to robust privacy”

¹ Le Morvan (2018) offers a way of dealing with *Mistaken Conclusion*: deny that privacy is factive (see also Allen (1988)). I set this suggestion aside here because it doesn’t accommodate *Incomplete Investigation*.

² Véliz’s preferred account of privacy differs slightly from EAP, see Véliz (2024: 79).

(Véliz 2024: 149). When you affect privacy (as characterized by EAP) you necessarily affect *robust* privacy, but the converse doesn't hold because robust privacy is affected if there are certain possible worlds in which privacy *would be* affected.

Véliz's account holds promise in explaining how to violate someone's right to privacy while leaving their privacy unaffected. She glosses the point as follows:

I violates S's right to privacy if he (...) secures a position from which he can invade S's privacy with the intention of invading S's privacy now, in the future, or in some counterfactual circumstance (...). (Véliz 2024: 124-125)

It is instructive to work through the cases with this in hand. Véliz's condition can be used to handle *Incomplete Investigation* because Barrack secures a position (by collecting the footage) from which he can now easily make Dough lose privacy if he so intended (by looking at the footage). In possible world talk, whether Dough loses privacy is counterfactually dependent on little but Barrack's intentions. Moreover, it seems that the condition can deal with *Mistaken Conclusion*. Even if Shoshana forms false beliefs in the actual world, she nevertheless secures a position from which she can easily learn facts about Linda. To see this, suppose that Shoshana decided to re-evaluate her (assumed to be excellent) evidence at a later point in time; not many things would have to change about the actual world for her to learn the truth.

There are at least two problems with Véliz's account. First, it's a strikingly revisionary response to the alignment challenge to say that, suitably interpreted, the right to privacy isn't strictly speaking about privacy but something rather different, *robust* privacy.³ I'll press a different worry here, though, suggesting that Véliz's solution comes at the cost of implausibly inflating the scope of the right to privacy. More specifically, Véliz's account doesn't steer clear of a challenge originally pressed against so-called control accounts of privacy famously due to Thomson ((1975: 305), see also Parent (1983) and Rickless (2007)):

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

³ Véliz offers some reasons for why this conclusion is acceptable. See for example Véliz (2024: 76, 143, 148)

Thomson is right, or so I shall assume.⁴ But it's not clear that an account of the right to privacy that is sensitive to what happens in possible worlds can accommodate it. To see this, consider an example inspired by Thomson's observation:

Acquiring The Means. Sophia buys an X-ray device. She buys it because she plans to spy on Miranda but doesn't follow through on her plan.

Sophia acquires the device with the intention of spying. In this case, Véliz's condition is satisfied because Sophia secures a position from which she can easily invade Miranda's privacy (all she must do is deploy the device). Intuitively, though, Sophia doesn't violate Miranda's right to privacy despite there now being a nearby possible world in which Sophia would reduce Miranda's privacy. Paraphrasing Thomson, Miranda's right to privacy is not violated until Sophia uses the device. Perhaps Véliz could respond by restricting what nearby possible worlds we should take in account. But since it seems that the important detail here is what happens in the actual world, it's not clear that this would work as desired. Véliz's solution to the alignment challenge, it seems, falls prey to Thomson's influential challenge.

3. The Inquiry Account

In laying out my preferred response, I'll begin with an observation. According to EAP, whether A is in a state of privacy is a function of whether B is in a certain epistemic state. But epistemic states – in virtue of being states – are not natural candidates for having the kind of normative statuses that we associate with violations of moral rights (but see Basu (2023a; 2023b)). This is one reason why it sounds about right when Thomson says, “You may violate a man's right to privacy by looking at him or listening to him; there is no such thing as violating a man's right to privacy by simply knowing something about him.” (Thomson 1975: 307).

This observation provides a rationale for thinking that what privacy is, on the one hand, and the activities that could potentially violate someone's right to privacy, on the other, couldn't exactly be overlapping. But it doesn't offer us an alternative. However, if EAP is on the right track, we can use it to develop an account of how these things might be related, by focusing on a

⁴The challenge I press against Véliz might seem question-begging because she endorses a very expansive account of the scope rights. But, as noted earlier, I think we should prefer a solution to the alignment challenge that avoids posing a non-standard account of the scope of rights. More specifically, by equating the act of *gaining the power to spy* with (actually) *spying* in terms of their privacy right-violating potential, it seems to me that we inappropriately flatten the normative landscape and lose sight of what is distinct about privacy in the first place.

distinctive kind of activity that bears a structural relation to epistemic states, namely the activity of *inquiry*.

We can develop a rough account of inquiry by drawing on recent work in epistemology (e.g., Friedman (2019); (2020); Thorstad (2021); (2022); Steglich-Petersen (2021); Fleisher (2023); Flores and Woodard (2023); Carter and Hawthorne (forthcoming); Haziza (2023)). For someone to count as being engaged in inquiry, two conditions must be satisfied. First, when engaged in inquiry you possess some mental state directed towards the goal of learning the answer to some question, something Friedman calls “interrogative attitudes” (Friedman 2019). Examples of such interrogative attitudes are *wondering if p*, *being curious if p*, or *deliberating if p*. This condition implies, for example, that if you are looking at the sky but without being directed by the goal of learning what the weather is like, you are not inquiring into what the weather is like (suppose you look up only to stretch your neck to alleviate pain). Second, you inquire into *p* only if you take certain steps towards this goal. To illustrate, you are not inquiring into *p* if you desire to know if *p* but do nothing to bring yourself closer to this goal. On the other hand, there are many different types of things you could be doing while guided by goal-directed attitude and count as inquiring. For example, inquiry may involve psychological processes such as weighing evidence, making inferences, entertaining alternative hypotheses, and reflecting upon the sufficiency of one’s available evidence. But non-psychological processes also form part of inquiries. For example, going to the library to search for books (gathering evidence) or having a conversation with a friend about whether it’ll rain tomorrow (shared deliberation) are activities that can form part of inquiries. Epistemologists have, however, tended to focus on evidence-gathering as the paradigmatic activity associated with inquiry (e.g., Carter and Hawthorne, forthcoming). Finally, it’s worth noting that inquiry is a temporally extended activity. Inquiries can be extended over long durations of time and involve many sub-activities that are strung together as subplans to achieve the goal of knowing the answer to a question (e.g., a research project). But they can also be brief (e.g., inquiring into whether it’s raining outside just now by looking outside the window).

From the characteristic of the activity of inquiry, we can observe that there seems to be a systematic relationship between inquiry and epistemic states, even if the exact details of this relationship might be hazy (see, e.g., Steglich-Petersen (2021) for discussion): Inquiry is *the* central intentional activity that rational agents engage in to achieve epistemic states such as knowledge, certainty, or rationally justified belief. This observation is crucial to solving the alignment challenge since, if privacy is a direct function of epistemic states, it would come as no surprise that acts that violate the right to privacy can be described as instances of inquiry, given the systematic,

instrumental connection there is between this activity and epistemic states. I take this observation to motivate:

The Inquiry Account of the Right to Privacy (IA): All acts that violate the right to privacy can be described as instances of inquiry.

Notice that IA isn't a complete account of the right to privacy. Rather, it supplies a necessary condition for when the right to privacy is violated. To solve the alignment challenge, we don't need a full account of the right to privacy. What we need is a systematic answer to how what is protected by the right to privacy could reasonably be taken to be about privacy (as analyzed by EAP). One way to illustrate this is by the observation that solving the alignment challenge shouldn't foreclose substantive normative disagreement about the adequate scope of the right to privacy. IA accommodates this by leaving room for normative considerations explaining why, and when, an act of inquiry violates the right to privacy.

Let's consider how IA solves the alignment challenge facing proponents of EAP. First, it provides the following rationale for how what privacy is and the right to privacy are connected: Even if privacy and the right to privacy concern different things (epistemic states and inquiry, respectively) there is nevertheless a deep, structural connection between them. This is because inquiry stands as the central activity directed towards the goal of forming epistemic states. When you inquire, you pursue the answer to a question. And, as Haziza (2023: 1) notes, answers are propositions. Knowledge of certain propositions, in turn, is the kind of things that undermine privacy according to EAP. The right to privacy, then, is a right against people engaging in (certain forms of) inquiry since avoiding this activity is what could prevent them from undermining other's privacy.⁵ In other words, the right to privacy and privacy as a phenomenon stand in an instrumental relationship of means and ends.

Second, IA delivers appropriate verdicts on the cases that plagued the other accounts we considered in the previous section. Consider first *Incomplete Investigation*. Barrack violates Dough's right to privacy because he inquires by way of collecting evidence. The fact that Barrack never comes around to complete this investigation (by looking at the footage) is immaterial since what he is doing can be described as an incomplete inquiry. Barrack shouldn't (without consent) inquire into the kind of affairs that can become known by collecting evidence via a camera installed in somebody's home. Since *Incomplete Investigation* is analogous to the case of the NSA that opened

⁵ Véliz notes that the right to privacy has to do with ways of getting at the object of privacy (Véliz 2024: 76, 143).

this paper, it can be dealt with similarly: Even if the NSA never looks at the information, they collected about you, they are nevertheless wrongfully inquiring, which explains why this could be violating your right to privacy.⁶

Consider next *Mistaken Conclusion*. The fact that Shoshana forms false beliefs implies that Linda's privacy is intact (following EAP). But Shoshana can nevertheless be said to be engaged in inquiry and may therefore be said to violate Linda's right to privacy. In this way, IA seems to deliver the desired verdict.

Next, IA drives a wedge between cases like *Mistaken Conclusion* and *Incomplete Investigation*, on the one hand, and cases like *Acquiring The Means* on the other. How so? Even if Sophia buys the X-ray device because she intends to use it to spy on Miranda, it sounds odd to say that Sophia is engaged in inquiry when buying the device. It seems much more natural to say that Sophia initiates her inquiry when she deploys the device to spy on Miranda. In other words, counting the acquisition of the X-ray device as a part of Sophia's inquiry into Miranda's personal affairs would conflate the activity of inquiring with something much broader: the activity of *acquiring the means* to inquire. We can observe this by focusing on the paradigmatic activity associated with inquiry, the activity of evidence-gathering: When Sophia acquires the X-ray device, she acquires the means to gather evidence about Miranda's private affairs, but it is only when she begins gathering evidence (by deploying the device) that inquiry is initiated. Accordingly, the IA implies, attractively, that Sophia is *not* violating Miranda's right to privacy.

In closing, I'd like to consider an objection to IA. Inquiry, as noted above, is an intentional or goal-directed activity. But it would seem possible to violate somebody's right to privacy without inquiring. Marmor (2015: 18) puts the point as follows,

Suppose, for example, that Bob (...) plays with his fancy new X-ray machine, aiming it at what he takes to be the bare wall in Mary's house. It is not the nicest thing to do, you might think, and perhaps it is even careless and negligent, but it is not a deliberate attempt to find out anything, either [...] I would be inclined to think that Bob violated Mary's right to privacy, inadvertently.

In a sense, this is a counterexample to the inquiry account because Bob violates Mary's right to privacy inadvertently, without being engaged in inquiry as I have characterized the activity. I don't

⁶ A complicating factor that I suppress here is that the NSA is a group agent. But groups can inquire and know as well (see Habgood-Coote 2022).

think this shows, however, that we should discard IA. Instead, it shows that we should amend it with terminology to describe the cases where you (wrongfully) engage in activities conducive to the goal of forming epistemic states, but without doing so under a goal-directed description. It's worth adding that there is nothing mysterious about this. For example, criminal law typically distinguishes between intentional species, on the one hand, and negligent or reckless species, on the other, of the same kind of wrongdoing. Based on that, we can for example distinguish manslaughter (intentionally killing a human) and homicide (negligently killing a human), but it's relatively clear that they belong to the same genus of wrongful conduct (that of killing a human).

From this perspective, the IA looks attractive because it helps us characterize the genus of the activities that violate privacy rights: it involves undertaking activities that would serve as the direct means to achieve the aims of a rational agent motivated by answering a certain question. Or simpler, taking steps that could form part of an inquiry. Interestingly, this picture aligns nicely with Thomson's rough description of activities that violate the right to privacy as activities that involve "taking certain steps to find out facts" (Thomson 1975: 307). In this way, even if IA is in one respect too restrictive, it is nevertheless explanatorily fruitful because the paradigmatic activity of inquiry can be used to explain in what sense identical activities that lack the mental states necessary for counting as inquiry could nevertheless count as violations of privacy rights. This strikes me as sufficient to solve the alignment challenge.

4. Conclusion

Proponents of the epistemic (access) account of privacy face the challenge of explaining why violations of the right to privacy that leave privacy unaffected are plausibly interpreted as violations of the right to privacy without casting doubt on their favored account. In this paper, I've offered a promising answer to this challenge in the form of the inquiry account. According to this, privacy is a function of epistemic states, and violations of the right to privacy (must therefore) range over a type of general activity we can describe as inquiry given the structural relationships between epistemic states and inquiry.⁷

Funding: This work was supported by the Carlsberg Foundation (grant number CF20-0257)

Author: Lauritz Aastrup Munch (lauritzmunch@gmail.com), Aarhus University (Denmark), Department of Philosophy and History of Ideas

⁷ **Acknowledgements:** I thank two referees from this journal for very helpful comments and suggestions.

References

- Allen, Anita L. 1988. *Uneasy Access: Privacy for Women in a Free Society*. New Jersey: Rowman & Littlefield Publishers.
- Basu, Rima. 2023a. Morality of Belief I: How Beliefs Wrong. *Philosophy Compass* (7):1-10.
- Basu, Rima. 2023b. Morality of Belief II: Three Challenges and An Extension. *Philosophy Compass* (7):1-9.
- Blaauw, Martijn. 2013. The Epistemic Account of Privacy. *Episteme* 10 (2): 167–77.
- Carter, Sam & Hawthorne, John. *Forthcoming*. Dogmatism & Inquiry. *Mind*.
- Fallis, Don. 2013. Privacy and Lack of Knowledge. *Episteme* 10 (2): 153–66.
- Fleisher, Will. 2023. Intellectual courage and inquisitive reasons. *Philosophical Studies* 180 (4):1343-1371.
- Flores, Carolina, and Elise Woodard. 2023. Epistemic Norms on Evidence-Gathering. *Philosophical Studies* 180 (9): 2547–71.
- Friedman, Jane. 2019. Inquiry and Belief. *Noûs* 53 (2): 296–315.
- Friedman, Jane. 2020. The Epistemic and the Zetetic. *Philosophical Review* 129 (4): 501–36.
- Habgood-Coote, Joshua. 2022. Group Inquiry. *Erkenntnis* 87 (3):1099-1123.
- Haziza, Eliran. 2023. Norms of Inquiry. *Philosophy Compass* 18 (12): 1-11.
- Le Morvan, Pierre. 2015. Privacy, Secrecy, Fact, and Falsehood. *Journal of Philosophical Research* 40: 313–36.
- Le Morvan, Pierre. 2018. Information, Privacy, and False Light. in Cudd, A., Navin, M. (eds) *Core Concepts and Contemporary Issues in Privacy*. Springer.
- Lundgren, Björn. 2020. A Dilemma for Privacy as Control. *Journal of Ethics* 24: 165-175.
- Lundgren, Björn. 2021. How We Can Make Sense of Control-Based Intuitions for Limited Access-Conceptions of the Right to Privacy. *Journal of Ethics and Social Philosophy* 20 (3): 382-391.
- Macnish, Kevin. 2018. Government Surveillance and Why Defining Privacy Matters in a Post-Snowden World. *Journal of Applied Philosophy* 35 (2): 417–32.
- Mainz, Jakob Thrane & Uhrenfeldt, Rasmus. 2020. Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy. *Res Publica* 27 (2): 287-302.
- Marmor, Andrei. 2015. What Is the Right to Privacy? *Philosophy and Public Affairs* 43 (1): 3–26.
- Matheson, David. 2007. Unknowableness and Informational Privacy. *Journal of Philosophical Research* 32: 251–67.

- 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.
- Menges, Leonhard. 2021. A Defense of Privacy as Control. *Journal of Ethics* 25: 385–402.
- Munch, Lauritz, and Jakob Mainz. 2023. To Believe or Not to Believe - That Is Not the (Only) Question: The Hybrid View of Privacy. *The Journal of Ethics* 27 (3): 245–61.
- Parent, William. A. 1983. Recent Work on the Concept of Privacy. *American Philosophical Quarterly* 20 (4): 341–55.
- Pettit, Philip. 1997. *Republicanism: a theory of freedom and government*. New York: Oxford University Press.
- Rickless, Samuel C. 2007. The Right to Privacy Unveiled. *San Diego Law Review* 44 (1): 773–99.
- Steglich-Petersen, Asbjørn. 2021. An Instrumentalist Unification of Zetetic and Epistemic Reasons. *Inquiry*: 1–28.
- Thomson, Judith Jarvis. 1975. The Right to Privacy. *Philosophy and Public Affairs* 4 (4): 295–314.
- Thorstad, David. 2021. Inquiry and the Epistemic. *Philosophical Studies* 178 (9): 2913–28.
- Thorstad, David. 2022. There Are No Epistemic Norms of Inquiry. *Synthese* 200 (5): 1–24.
- Véliz, Carissa. 2024. *The Ethics of Privacy and Surveillance*. Oxford: Oxford University Press.