




ARTICLE

Can privacy be diminished by falsehoods?

Alice Schneider 

Stanford Law School, Stanford University, Stanford, CA, USA
Email: alice898@law.stanford.edu

(Received 18 November 2024; revised 4 January 2025; accepted 17 January 2025)

Abstract

It is widely presumed that privacy is ‘factive’, i.e. that it cannot be diminished by accessing or disseminating falsehoods. But if this is so, what wrongs are committed in cases where others access documents of ours (letters, medical records, etc.) which contain false information? In this article, I examine various ways of explaining the wrongfulness of accessing and dissemination falsehoods (defamation; that privacy can be violated without being diminished; ‘control’ accounts of privacy; downstream revelations of truths; that falsehoods diminish ‘propositional’ or ‘attentional’ privacy). I lay out what each of these accounts misses about accessing falsehoods, about privacy, and/or about the right to privacy. I then propose two alternative ways of accounting for the intuitive wrongfulness of accessing and disseminating falsehoods: viewing them as merely ‘attempted’ privacy violations and weakening the truth condition of privacy diminishment.

Keywords: Privacy; falsehoods; defamation; factive

1. Introduction: accessing falsehoods

Access-based accounts of privacy generally define a loss of privacy as involving a change in a ‘relational epistemic state’¹, or as an ‘epistemic event’². The central case epistemic state that is thought to diminish privacy is thought to be that of *knowledge*; someone loses privacy when we *know* sensitive personal information about them. In the words of Klemens Kappel, ‘privacy depends inversely on epistemic access’³. Similarly, David Matheson explains that ‘[a]n individual A has informational privacy relative to another individual B and to a personal fact about A if and only if B does not *know* f.’⁴ Some, like Martijn Blaauw and Carissa Véliz, have proposed to weaken the justification requirement and to view that ‘weak knowledge’, like lucky guesses, might still diminish

My thanks for comments and extensive discussion go to Davin Lunz, Marc Schneider and Lynn Schneider.

¹Martijn Blaauw, ‘The Epistemic Account of Privacy’ (2013) 10 *Episteme* 167, 167.

²Carissa Véliz, *The Ethics of Privacy and Surveillance* (Oxford University Press 2024) 162.

³Klemens Kappel, ‘Epistemological Dimensions of Informational Privacy’ (2013) 10 *Episteme* 179, 185.

⁴David Matheson, ‘Unknowableness and Informational Privacy’ (2007) 32 *Journal of Philosophical Research* 251, 259. (Emphasis added).

privacy.⁵ Putting the justification element aside, theorists seem overwhelmingly committed to the truth requirement. They view that only *true* beliefs are capable of diminishing privacy⁶ and that privacy is ‘factive’⁷. Martijn Blaauw seeks to illustrate this in the following hypothetical:

Murderer 1

John believes his father to be a convicted murderer. His mother always told him so. As it turns out, however, his father wasn’t the convicted murderer John thought him to be at all. His mother merely misled him into thinking this, for reasons of her own. John has never told anyone that his father was a convicted murderer because he was embarrassed and wanted to keep this information⁸ private.⁹

Blaauw is adamant on the point that John cannot have privacy regarding the proposition that his father was a murderer. Given that John was mistaken about the truth of the proposition, he must also have been mistaken about its privacy: ‘John merely *thought* that he had privacy with respect to this false proposition. But in the end, there was nothing to be private about.’¹⁰ The case of *Murderer* is animated by the sense that it is simply logically impossible, or, in Blaauw’s own words, *incoherent*¹¹, to be said to have privacy regarding falsehoods.

Pace Blaauw, I will suggest below that our intuitions on this point are not quite so resolute. Consider the following sequels to the *Murderer* case:

Murderer 2. When visiting John’s home, I come across a letter from John’s mother, addressed to John, which discusses his supposedly murderous father. I open the letter and read it, coming to believe that John’s father was a convicted murderer.

Murderer 3. Suppose after reading the letter, I tell all my friends that John’s father is a convicted murderer.

On Blaauw’s view, the (intuitively transgressive) acts of, first, my reading the letter (*Murder 2*), and, second, my telling others about it (*Murderer 3*), do not diminish John’s privacy because to say that I have diminished privacy by accessing or disseminating falsehoods is a category mistake. Blaauw tells us ‘[t]o expose someone’s privacy is to reveal personal information about them, and one can only reveal facts, not falsehoods’.

We can expect that John would initially feel his privacy diminished upon learning that I read his mother’s letter and spread the news. Yet, in favour of Blaauw’s view, we might grant that John is liable to change his mind if he were to subsequently learn that his father was not a convicted murderer after all. We can imagine John’s relief in

⁵Blaauw proposes that “for a subject S to have ‘full privacy’ about a particular, true proposition, certain other individuals are not to stand in any epistemic relation to this true proposition about S” Blaauw (n 1) 167. See also Véliz (n 2) 234. For a discussion of beliefs that fall below the threshold of justification required for knowledge, see Don Fallis, ‘Privacy and Lack of Knowledge’ (2013) 10 *Episteme* 153, 157.

⁶See, e.g., Véliz (n 2) 214; Blaauw (n 1) 169; Two exceptions to this are Kappel (n 3) 180. and Pierre Le Morvan, ‘Privacy, Secrecy, Fact, and Falsehood’ (2015) 40 *Journal of Philosophical Research* 313. I discuss their accounts at length below.

⁷Blaauw (n 1) 169; Le Morvan (n 6) 322.

⁸It may be objected that falsehoods cannot count as ‘information’ See Luciano Floridi, ‘In Defence of the Veridical Nature of Semantic Information’ (2007) 3 *European Journal of Analytic Philosophy* 31, 31. I comment on my (and Blaauw’s) use of the term ‘information’ towards the end of this introduction.

⁹Blaauw (n 1) 169.

¹⁰*ibid.* (emphasis added).

¹¹*ibid.*

exclaiming: ‘The information which was spread was not true! Phew! No private information was revealed after all; no harm done!’ This suggests that John’s initial sense of a privacy diminishment (or even violation) might indeed have been mistaken.

On the other hand, it is equally plausible that John might not be so generous. (We can imagine him saying: ‘So what if it turns out that what the letter said isn’t true after all. You read my private letter without my permission! And then you then told everyone about it! The harm is done!’) Even if the information about his father turns out to be false, my accessing and spreading of it likely created just as much embarrassment, reputational harm and disadvantage for John as it would have, had it been true.¹² It seems undeniable that even after learning that the information I revealed was false, John would not be wrong to remain upset with me. Intuitively, in reading his letter (*Murderer 2*) and disseminating the information in it (*Murderer 3*), I have committed *some kind of wrong* against John¹³. The most intuitive way of explaining this wrong would be to say that I have diminished John’s privacy in a way that is somehow illicit, ergo I have *violated* John’s privacy¹⁴. But the view that privacy is factive cuts into this: if privacy cannot be diminished by falsehoods, then what is the wrong I have committed?

My argument here proceeds from the assumption that the acts of reading the letter John’s mother wrote to him (*Murderer 2*) and the subsequent dissemination of this information (*Murderer 3*) are somehow wrongful. In what follows, I lay out various ways of explaining this wrongfulness. Firstly, I review various ways of accounting for the wrongfulness of *Murderer 2* and *3* that have been offered in, or can be deduced from, the academic literature. Specifically, I consider defamation (section 2), that privacy can be violated without being diminished (section 3), accounts of privacy as control over information (section 4), downstream revelations of truths (section 5), and that falsehoods diminish ‘propositional’ or ‘attentional’ privacy (section 6). I reject these approaches as either being unable to account for what is intuitively wrongful in *Murderer 2* and *Murderer 3*, or as offering an unattractively broad account of privacy (or the right to privacy). Then, in sections 7 and 8, I propose two alternative ways of accounting for the intuitive wrongfulness of accessing and disseminating falsehoods: inchoate offences (i.e. attempted privacy violations; section 7), or weakening the truth condition of privacy diminishments (section 8).

Before I begin, I should clarify out that I use the term ‘information’ in a generic, non-factive sense. In doing so, I am following Blaauw, who himself uses the term ‘information’ to refer to the falsehood that John’s father is a murderer (I quote his hypothetical *Murderer* above). It might be objected that information is factive¹⁵, and that the statement ‘John’s father is a convicted murderer’ cannot count as ‘information’ if false. If one assumes that ‘information’ is the proper object of privacy, one might think that if information is itself factive, then privacy must be factive by extension. Ergo, no privacy diminishments by falsehoods.

This, I submit, is questionable. First, the question of whether information is factive is far from settled¹⁶. Even those who view that information is factive (or ‘veridical’) grant that, in

¹²In a similar case that I will discuss further below, Klemens Kappel describes the resultant harm as “a harm that is in all respects similar to a privacy harm, except that it derives from a falsehood.” Kappel (n 3) 190.

¹³One could deny that accessing falsehoods is wrong at all. On that view, no wrong has been committed in the cases of *Murderer 2* and *3*. My argument proceeds from the premise – which I take to be more intuitively compelling – that the conduct in *Murderer 2* and *3* is wrong *in some sense*.

¹⁴I will comment more on the distinction between normatively neutral ‘privacy diminishments’ and normatively laden ‘privacy violations’ – which are wrongs – below.

¹⁵I thank my reviewer for pointing this out.

¹⁶I address this debate in (slightly) more detail in section 8 of this article.

ordinary language, ‘information’ can be used as a synecdoche to refer both to ‘information’ and to ‘misinformation’¹⁷. This is the more generic sense in which I (and Blaauw) use the term. Second, and more importantly, it is plausible that the set of things covered by privacy is not limited to true information (to simply assert so would be to beg the question), but that it includes things, including (as I will suggest in section 8) certain kinds of falsehoods. In short, whether privacy is factive does not boil down to whether information is factive.

Aside from how I use the term ‘information’ in this article, it is worth pointing out that discussions of whether information is factive are analogically relevant to my argument about privacy. The question of whether ‘information’ is factive – much like whether verbs like ‘learning’, ‘remembering’, ‘revealing’ and ‘realizing’ are factive – remains a matter of continuing disagreement amongst philosophers. This genre of disagreement exposes contrasting intuitions which mirror differing intuitions on whether privacy is factive. As I will propose in section 8, these intuitive disagreements may be attributed to a difference in methodology within conceptual analysis.

2. Are cases of accessing and spreading falsehoods defamation?

As many have been quick to point out¹⁸, the canonical wrong of propagating falsehoods is defamation, i.e. libel or slander. Propagating the false claim that someone is a murderer – as I did in *Murderer 3* – seems like a classic case of defamation¹⁹. So we might view that in telling all my friends about John’s supposedly murderous father, I may have *thought* I committed the wrong of a privacy invasion (given I assumed the information about John’s father was true), but the wrong I effectively *committed* – somewhat ironically – was slander.²⁰

The problem with this view is that defamation does not neatly capture all intrusions that would, if relevant information were true, qualify as privacy invasions. Firstly, legal actions in defamation typically require demonstrable economic or reputational harm. This is different for privacy invasions, which are viewed as completed as soon as relevant personal information is *accessed*,²¹ irrespective of whether any downstream reputational or economic harm ensues. A further distinction is that actions in defamation typically require that the defendant (or, in civil actions, the plaintiff) acted with the knowledge that the information they published was false, or at least that they were reckless in deciding to publish the statement without investigating its accuracy.²² In other words, in a case such as *Murderer 3*, where I reasonably (and sincerely) took the information about John’s father to be true, would be ruled out as defamation. Privacy, at least in law, protects more tightly.

¹⁷Luciano Floridi, *The Philosophy of Information*, vol 9 (OUP 2011) 104. For a full defence of an alethically neutral conception of ‘information’, see Björn Lundgren, ‘Does Semantic Information Need to Be Truthful?’ (2019) 196 *Synthese* 2885.

¹⁸Parent tells us: “The spreading of falsehoods or purely subjective opinions about a person does not constitute an invasion of his privacy. It is condemnable in the language of libel or slander.” WA Parent, ‘Privacy, Morality, and the Law’ (1983) 12 *Philosophy and Public Affairs* 269. See also Véliz (n 2) 217.

¹⁹See, e.g. *Condit v. Dunne*, 317 F. Supp. 2d 344 (S.D.N.Y. 2004), as cited by Golden, T.H. and Boyle, C. (2013). *United States*. In *International Libel and Privacy Handbook*, C.J. Glasser (Ed.), p. 105.

²⁰Correspondingly, we might view that someone who seeks to defame another person by making up the vicious lie that their father is a murderer ends up violating the victim’s privacy if it turns out she providentially told the truth, see Véliz (n 2) 234. Fallis would disagree in cases where the (true) belief is not in any way ‘hooked up’ to the fact of John’s father being a murderer; Fallis (n 5) 157.

²¹Legal protections of privacy typically hook on to antecedent actions like collecting, transferring and disseminating information. This is sensible, given that doxastic involuntarism (for the view that privacy sometimes requires us to suspend our beliefs, see Munch (n 61) 553.

²²*New York Times Co. v. Sullivan*, 376 U.S. 254 (1964).

One might object that these distinctions are merely a matter of positive law, but not of moral merit. That legal actions in defamation claims in systems like the US and the UK happen to have these various prerequisites should not bar us from using the *moral* concept of defamation to account for cases of propagating falsehoods (i.e. the case I describe as *Murderer 3*). So, even though legal actions of defamation might fail in the courts, we might view that *morally speaking*, in the case of *Murderer 3*, I committed the wrong of defamation.

But plugging the (moral) wrong of defamation in cases where the person accessing relevant information clearly took themselves to be spreading *true* information seems odd. Defamation and privacy violations may be complimentary wrongs in the sense that they are both potentially concerned with reputational harm. But their respective relationships with truth have opposing normative implications: the kinds of beliefs about another person that are thought to diminish privacy typically require some inquiry work – some digging or searching or hacking – or they involve exploiting that one already stands in an epistemically privileged (i.e. confidential) relationship to the victim. As Kappel argues, what makes privacy violations intuitively wrongful has to do with the ‘epistemic pathway’, the *how* one comes by this information.²³ The wrongness of defamation seems inverted: defamatory falsehoods are exceedingly cheap to come by; they can just be made up, out of thin air. Indeed, any relevant epistemic effort (i.e. digging or searching or hacking) in the context of coming by information used for defamatory purposes does not ground wrongdoing, but instead counts in the perpetrator’s *favour* insofar it suggests that active steps were taken to verify the claims’ accuracy²⁴. Put crudely, privacy invasions are about checking propositions about others too thoroughly, while defamation is, at best, about not checking propositions thoroughly enough (and, at worst, about intentionally propagating what which one knows to be false). In defamation cases, truth counts as a defence; it is what licenses spreading statements which cause reputational and economic harm. The normative vectors of the wrongs of defamation and privacy invasions clearly point into opposing directions.

There are many ways of defaming someone that do not at all involve sharing information which is private or personal, or which *would be* private had the information been true. The International Libel and Privacy Handbook lists the following classic examples for statements found to be defamatory under U.S. law: ‘playing a role in an alleged kidnapping and murder’²⁵; ‘incompetence of a professional’²⁶; or ‘being a communist’²⁷. More recently, the far-right radio show host Alex Jones repeatedly claimed that the 2012 Sandy Hook Elementary School shooting was a hoax and that the families of the victims were lying about the tragedies they suffered (the families of victims faced harassment and death threats as a consequence). Jones was held liable to pay roughly 960 million USD in damages for defamation.²⁸ In none of these cases would the communicated falsehoods qualify as protected by privacy, had they been true.

²³ “[C]ertain pathways are illegitimate, while other ways of gaining access to the very same facts are not”. Kappel (n 4) 190. Judith Jarvis Thompson takes the view that “simply knowing” cannot constitute a privacy violation; the violation must depend on how the information was come by. Judith Jarvis Thompson, ‘The Right to Privacy’ (1975) 4 *Philosophy and Public Affairs* 295, 4.

²⁴ It is of course possible that someone engages in digging or searching or hacking (i.e. invades another person’s privacy) to find information that is then be marshalled to construe particularly hurtful defamations.

²⁵ *Condit v. Dunne*, 317 F. Supp. 2d 344 (S.D.N.Y. 2004), as cited by Golden, T.H. and Boyle, C. (2013). United States. In *International Libel and Privacy Handbook*, C.J. Glasser (Ed.), p. 105.

²⁶ *Scripps Texas Newspapers v. Belalcazar*, 99 S.W.3d 829 (Tex. 2003), as cited by Golden, T.H. and Boyle, C. (2013). United States. In *International Libel and Privacy Handbook*, C.J. Glasser (Ed.), p. 105.

²⁷ *MacLeod v. Tribune Publ’g Co., Inc.*, 52 Cal 2nd 536 (Cal. 1959) as cited by Golden, T.H. and Boyle, C. (2013). United States. In *International Libel and Privacy Handbook*, C.J. Glasser (Ed.), p. 105.

²⁸ *Lafferty v. Jones*, 336 Conn. 332, 246 A.3d 429 (Conn. 2020).

On the side of privacy, we can equally conceive of many cases of spreading private information which would not count as defamation if false, because they would not be ruinous to one's reputation. Consider the case of a proud mother who violates her adult son's privacy by sharing his excellent GRE results on her public Facebook page, even though she knows that the son does not want this information to become public. If the son's GRE results turned out to be false due to some system error, what the mother did would still not count as libel or slander; the sharing of the test results is aimed at promoting, rather than diminishing, her son's reputation.

For the purposes of my argument here, a final, compelling reason to remain unsatisfied with the view that defamation can mop up cases falsehoods is that the wrong of defamation is incapable of accounting for the intuitive wrongfulness *Murderer 2*, i.e. the variant where I read (i.e. access) the letter John received from his mother and keep the contents to myself. Defamation is a wrong of an inherently *communicative* nature, so we can hardly stretch it to cases of 'mere' access.

3. Can we commit privacy violations without diminishing privacy?

The intuitive wrongfulness of accessing falsehoods (i.e. *Murderer 2*) has worried other access theorists. Notably, Klemens Kappel suggests we allow that privacy diminishment and privacy violations (or 'wrongs', as he calls them), come apart in instances where someone forms a false belief based on illegitimately accessing sensitive information about another person. In Kappel's hypothetical (which I will proceed to call *Adam*), I break into my neighbour Adam's house and read a letter from which I mistakenly infer (because of my poor German) that Adam is a member of a neo-Nazi party (in fact, the letter contains a rejection of an application for party membership)²⁹. Kappel contends that Adam's privacy has not been diminished by my reading the letter, given that the belief I form as a consequence is false. Nevertheless, he maintains, our intuitions mandate that this case must still count as a privacy violation.³⁰

The distinction between privacy *simpliciter* and a normatively laden *right to privacy* is fairly common in the philosophical literature³¹, particularly amongst access theorists who view – somewhat controversially, in my view – that privacy *simpliciter* is lost or 'diminished' every time personal information is divulged voluntarily to another person³², but who do not want to normatively label such innocuous cases privacy violations. On this view, not all privacy diminishment are privacy violations, so the *right to privacy* is drawn more narrowly and thus only covers a strict subset of the things that are part of privacy *simpliciter*.

However, what Kappel suggests when he argues that it is possible to violate a right to privacy – without *actually diminishing* privacy – is that the sets of things covered by privacy *simpliciter* and by the *right to privacy* merely intersect. In other words, he, in some respect, draws the right to privacy more *broadly* than privacy *simpliciter*. This begs the question of what exactly this supposed *right to privacy* is tracking. Lauritz Munch

²⁹Adam is in fact an under-cover anti-Fascist who seeks to expose the party. (Yes, it's complicated!) See Kappel (n 3) 190.

³⁰See *ibid.* Véliz equally argues that though the attention prompted by false beliefs can risk privacy (Véliz (n 2) 218.), privacy is not diminished by falsehoods.

³¹See, e.g. Véliz (n 2) 213.

³²By contrast, some control theorists view the voluntary sharing of information as an exercise of control over information and thus do not count it as a privacy diminishment. See, e.g. Julie C Inness, *Privacy, Intimacy and Isolation* (Oxford University Press 1992) 46: "Our impulse in these cases is to say that we are including another within our realm of privacy, not lessening our privacy".

and Jakob Mainz have described Kappel's way of distinguishing a right to privacy from privacy simpliciter as incurring a theoretical debt, since 'on any plausible account of the right to privacy, privacy must be the object of this right'.³³

One account that offers a more substantive explanation for how violations of privacy can be drawn more broadly than privacy diminutions is Carissa Véliz' 'hybrid account' of privacy. She distinguishes between privacy, which she defines roughly along the lines of access-theoretical accounts, and broader 'right to privacy' or 'robust privacy', which tracks other accounts of privacy as control over information³⁴. Accordingly, cases where 'A secures a position from which to invade S's privacy with the intention of invading her privacy [...] in some counterfactual circumstance'³⁵ already qualify as violations.³⁶ In other words, the *right to privacy* is violated even in cases where privacy simpliciter remains undiminished ('unaccessed'), such as a case where a friend acquires control of one's diary 'just in case she might feel like reading it in the future'³⁷. This is possible because Véliz defines the 'right to privacy' by adopting Philip Pettit's notion of 'robustly demanding goods'.³⁸ She argues that to be said to have 'robust privacy'³⁹, it is not enough that one's privacy remains free from violation under the narrow conditions of the here and now, but that privacy remains 'unaccessed' in alternative possible worlds: 'Robustly demanding goods are ones that require counterfactual assurances'.⁴⁰ What does it mean to assure something counterfactually? Pettit himself gives the example of love and explains:

If you love me, so the lesson goes, [...] You must also feel and offer me care independently of how I currently look, what I currently do, or how I am currently called. Shakespeare already made the point in Sonnet 116: 'Love is not love, Which alters when it alteration finds'.⁴¹

The contention is that if my love for you is built on exceedingly shallow, contingent criteria – that your name happens to be Ernest⁴², for example – then this 'love' is so thin that it does not properly count as love at all⁴³. Love, according to Pettit, is a robustly demanding good in that it requires assurances beyond the contingent world we happen to live in. Is this also true of privacy? Véliz certainly thinks so; if the only reason I am not reading my roommate's diary right now is that we presently get along amicably – but if we were to fall out, I *would* read it immediately – then I cannot be said to be respecting

³³Lauritz Munch and Jakob Mainz, 'To Believe, or Not to Believe – That Is Not the (Only) Question: The Hybrid View of Privacy' (2023) 27 *Journal of Ethics* 245, 258.

³⁴Véliz (n 2) 189. I include a more general discussion of how control-views deal with falsehoods below in section 4.

³⁵*ibid* 353.P.

³⁶Assuming that they are not justified.

³⁷'Even if my diary is written in code, such that she could not gain access to its content if she wanted to, it seems like I have a privacy claim against my friend that access theories cannot capture' Véliz (n 2) 186.

³⁸Philip Pettit, *The Robust Demands of the Good* (Oxford University Press 2015).

³⁹Véliz (n 2) 190.

⁴⁰*ibid* 189.

⁴¹Pettit (n 38) 12.

⁴²Pettit cites Gwendolyn's love for Jack in Oscar Wilde's comedy "The Importance of being Ernest" as an example for the kind of thin, overly conditional love that we cannot really count as love at all, see Pettit p 11.

⁴³Pettit's argument gets much mileage out of our intuitive repulsion at a love that depends on superficial, unfair conditions. But (it seems to me) that even if I were loved by another person for a contingent reason (beauty, fame . . .) I would – perhaps luckily – still be truly loved in *this world*, even though I grant the feeling is not very counterfactually stable.

my roommate's 'robust privacy'. In this way, Véliz' argument richly illuminates what it means to positively be guided by another person's privacy in genuine, non-instrumental ways.

Before I discuss this account further, I should point out that Véliz does not herself claim that her hybrid account of privacy protects falsehoods from access or dissemination. But we might read her account such that it accommodates the view that accessing falsehoods count as violations of the 'right to privacy' (without actually diminishing privacy *simpliciter*). One way to do this is to sketch the *truth* of the relevant information as part of the counterfactual world; we might thus contend that if you are disposed to reading my letter in a possible world where the content of the letter *was* true, you do not respect my *robust privacy*. Again, the violation here depends on what you *would do* if things were (slightly) different, so the charge regarding *Murderer 2* is not that you *did* in fact read the letter containing falsehoods, but that you *would (also) have read* it in a possible world where the information it contained was true.⁴⁴ (In *Murderer 2* and 3, the perpetrator who read the letter assumed the information was true, so we can grant that had it actually been true, they would have acted just the same.)

For the case of *Murderer 2*, this solution has some appeal because it captures the detail the perpetrator clearly demonstrated a lack of appreciation for the victim's privacy (even if we think they did not succeed in diminishing it). This view takes the perpetrator to task for the possible world they presumed themselves to be acting in (rather than blaming them for something they did not take themselves to be doing, i.e. defamation).

However, counting what one might have done in a counter-factual world as a *violation of a right* seems overzealous. Is 'failing to respect' another person's robust privacy akin to *violating her right to privacy*?⁴⁵ Véliz commits herself to the view that I am actively *violating* my roommate's *right to privacy* in the here and now – despite the fact that I am not even reading her diary – just because I *would* read it, if circumstances were slightly different (i.e. if we got into a fight)⁴⁶. That seems inflationary. It is questionable that the terminology of 'right' serves Véliz' argument well. Pettit himself characterises robustly demanding goods like love and friendship as morally valuable objectives which we ought to, *prima facie*, *pursue*. But he comes up short of suggesting that these goods coagulate into definable 'rights'. In other words, when I fail to love others in robust ways, I am doing *just that* – failing to love; failing to maximise some 'rich good'⁴⁷, acting in morally suboptimal ways – but I am not quite violating anyone's *right* to be loved⁴⁸. Rights, on a more orthodox philosophical understanding, are more modest deontological concepts than virtues or consequentialist ideals; they leave room for supererogation⁴⁹. Even if I lack the disposition to respect your privacy

⁴⁴I do not know if this is a version of her account of robust privacy that Véliz herself would subscribe to.

⁴⁵When it comes to talk about rights, we tend to view that others are "respecting our rights" as long as they are not violating them. "Respect for rights" thus has a thin meaning – one roughly analogous to conformity – and it does not require proactive appreciation in the sense of *respectfulness* or endorsement. On Véliz' view, by contrast, 'respect for robust privacy' requires a kind of proactive valuing. If this is not given, relevant situations are counted as infringements or violations. See Véliz (n 2) 353.

⁴⁶Pettit himself grants that (robust) love need not be unconditional, or granted to "extreme possibility" (Pettit 15). This means we have license to imagine some possible world where my love for my partner might cease, but it still counts as 'robust love'. Similarly, for Véliz, there ought to be some removed possible worlds which would not undermine respect for robust privacy.

⁴⁷Pettit (n 38) 6.

⁴⁸That we are entitled to demand relevant performances from each other depends, according to Pettit, on prior relationships, i.e. attachments p 40.

⁴⁹Heyd, David, "Supererogation", *The Stanford Encyclopedia of Philosophy* (Spring 2024 Edition), Edward N. Zalta & Uri Nodelman (eds.) <<https://plato.stanford.edu/archives/spr2024/entries/supererogation/>>.

counterfactually, across various other possible worlds, it seems intuitive that as long as I am not reading your diary in *this one* – however reluctantly or contingently – then I am not quite *in violation* of your right to privacy (at least not *yet*⁵⁰).

Accounts which draw the right to privacy more broadly than privacy itself are potentially able to count *Murderer 2* and *3* as privacy violations whilst holding that privacy (or privacy *simpliciter*) remains undiminished. But accounts which have been offered in the literature so far either fail to explain what, aside from privacy *simpliciter*, the (broad) right to privacy protects, or they draw what this right protects so broadly that they end up with an indefensibly broad view of what it is to *violate* the right to privacy.

4. Can control views of privacy account for the wrongfulness of accessing falsehoods?

‘Control views’ of privacy typically confirm violations of the right to privacy when another person establishes control over a person’s information without their permission; even when that information remains ‘unaccessed’⁵¹ (i.e. before beliefs with any relevant propositional content were formed).⁵² Control accounts thus view that a privacy violation has occurred at an earlier logistical step than access accounts. Some consider this a feature and some a bug: those who seek to describe widespread data collection practices of, e.g. the NSA as constituting a *loss* of privacy (rather than a mere privacy *risk*) can use control accounts to do so. Others take the view that control accounts are overinclusive in picking out privacy diminishment, especially in cases where relevant information technically comes under another person’s control, but never accessed. For example, against control accounts, Kevin Macnish has defended the view if I leave my diary on a table in a coffee shop and when I come back to retrieve it after 30 minutes, a stranger finds it and returns it to me – without ever having opened it or having intended to open it – then I have lost no privacy.⁵³

The merits and demerits of control views are beyond the scope of this article. But it is noteworthy that control views draw distinctions when it comes to false beliefs that access accounts of privacy pass over. Note that the case of Kappel’s *Adam* differs slightly from the *Murderer* cases: in the case of *Adam*, the letter from the Neo-Nazi party does contain sensitive, *true* information; the intruder just did not end up forming true beliefs about Adam’s Neo-Nazi party membership because of some epistemic failing on their part (i.e. lack of linguistic competence). By contrast, in the case of *Murderer 2*, the letter John’s mother wrote to John contained *false information*. For access accounts of privacy, it is irrelevant whether the relevant false beliefs are due to some epistemic failure, or whether they are due to information being false. *Any* failure to form true beliefs about another person will mean the other person’s privacy remains, to that extent, undiminished. *Why* the belief ends up being false – i.e. because the information is encrypted; because it is written in another language; because the information is false; because the information is simply misunderstood – does not matter.

By contrast, control accounts would confirm that a diminishment of privacy has occurred in cases where control over *true* information was established, even if it did not

⁵⁰I will discuss attempted violations of privacy in section 8 of this article.

⁵¹Véliz (n 2) 162.

⁵²See Kevin Macnish, ‘Government Surveillance and Why Defining Privacy Matters in a Post-Snowden World’ (2018) 35 *Journal of Applied Philosophy* 417, 420.

⁵³*ibid*; Björn Lundgren, ‘A Dilemma for Privacy as Control’ (2020) 24 *Journal of Ethics* 165, 169. Menges calls this the ‘threatened loss objection’ Leonhard Menges, ‘A Defense of Privacy as Control’ (2021) 25 *Journal of Ethics* 385, 386.

lead to true beliefs (i.e. in cases where information turns out to be written in a foreign language or encrypted; in cases where relevant information ends up being misunderstood). But control accounts would deny – at least to the extent that they view privacy as *factive* – that a diminishment of privacy has occurred in cases where relevant information was itself false (i.e. *Murderer 2 and 3*).

Is there any normative justification for treating these cases differently? We might grant that in cases where control views confirm privacy losses, *true* information was at stake. This would matter to those who define privacy as *factive* (i.e. as being concerned with truths). Then again, this view is begging the question when it comes to the wrongfulness of accessing falsehoods.

On the other hand, taking up a more pragmatic perspective, we might not think it matters much at which step in the belief formation we locate the epistemic shortcoming. After all, the line between encrypted information and false information can seem arbitrary. Imagine that, as part of an encryption mechanism, I insert ‘it is not true that’ before every sentence in my private correspondence, such as before the sentence ‘I have an intimate relationship with Anna’. What would the control view say about a person who acquires my letters and reads them, but who does not succeed to decrypt my writing and thus forms the false belief of ‘it is not true that she has an intimate relationship with Anna’? We might view that the (post-encrypted) information printed in the letter was false, in which case control accounts would tell us no control over *true information* was ever established (ergo no privacy loss). Or did the person who read my letters acquire control over *true information* (plus a layer of encryption)? In that case, control accounts would view that a privacy loss *did* occur.

Similarly, consider cases of linguistic ambiguity. Suppose that I, a German speaker who has lived abroad for many years, write my diary using a random mix of German and English. On one page, I write:

Heute früh habe ich das gift für meinen Mann gekauft. Ich werde es ihm heute Abend geben.

(This translates to ‘This morning, I bought the gift for my husband. I will give it to him tonight’.) When I write the word *gift*, I use it with its English meaning; i.e. ‘present’. But in German, ‘Gift’ means ‘poison’. Assume now that a German speaker acquires control of this page of my diary without my permission. Not realizing that I sometimes use random English words in my writing, they form the false belief that I am about to poison my husband. Is this false belief due to the information on the page being false, or due to some hermeneutic shortcoming on the intruder’s part? What I wrote is subjectively true (i.e. true if interpreted as pragmatically enriched by the idiosyncratic way I use languages⁵⁴; I *did* buy a present for him after all.), so control accounts would view that privacy was lost. But what I wrote is also *literally* false, ergo no privacy loss. Which view to take? The point is that there are various cases where seems arbitrary to view that a privacy loss turns on the whether the information was encrypted or false.

The bottom line is that control views may be more accommodating of the view that privacy losses have occurred in *some* cases where false beliefs are concerned. But their generosity is arbitrary. Therefore, they are not obviously preferable when it comes to explaining whether or not falsehoods diminish privacy. In any case, my objective here is to explain the intuitive wrongfulness of *Murderer 2 and 3*; i.e. of cases where *false*

⁵⁴Perhaps the lack of capitalization should have been a give-away. Or it may have been read as a sign that I am bad at grammar.

information was accessed. On control views, no privacy loss/ diminishment occurs in these cases, so they do not explain what, if anything, was wrongful in these cases.

5. Do falsehoods diminish privacy indirectly by pressuring the disclosure of truths?

In the academic literature on falsehoods, one frequently discussed insight is that propagating falsehoods may lead to privacy diminishments down the line. Ruth Gavison explains that falsehoods, if sufficiently spectacular⁵⁵, lead to a loss of anonymity and increased attention:

Even if the defamatory information is false, it attracts attention to the person in ways that may involve loss of privacy.⁵⁶

Gavison suggests that the circumstances of such enhanced scrutiny constitute a loss of ‘attentional privacy’⁵⁷, ‘or at least the threat of such a loss’⁵⁸. Similarly, Véliz points out that false information may create pressure on those it concerns ‘to reveal more than they would’ve otherwise wanted to’.⁵⁹ In other words, private facts may end up being divulged in defence of false rumours.⁶⁰ In an example Véliz gives, a person who is the victim of rumours about her sex life feels compelled to defend herself by disclosing that she is ‘single and celibate’.⁶¹

Applied to the *Murderer 2* and *3* cases, this view suggests that privacy is ultimately lost only when John feels pressure to fight these false claims about his father by divulging true pieces of personal information about himself. We can imagine that John might react to other people’s false beliefs about his father by telling the story of how his mother lied to him in her letters. A false rumour leading to some dark family secrets being revealed; this seems much like the characteristic ‘downstream’⁶² privacy loss Gavison and Véliz have in mind.

But it seems far from given that any divulgence of true facts will follow from accessing and spreading falsehoods in the *Murderer* cases. First, John might never learn that the claims about his father’s criminal status were false, ergo he will feel no pressure to divulge sensitive truths. Second, in the case of *Murderer 2*, John might never find out that another person read his letter and now holds the relevant false beliefs about his father’s criminal status. Again, there is no reason he would feel tempted to divulge truths. Third, even if John did find out the claims about his father’s criminal status were false and that others (falsely) believe them, he might just choose not to correct these false beliefs. Alternatively, John might choose to undermine others’ false beliefs by spreading yet another falsehood about his father (i.e. a falsehood he is more comfortable with, such as ‘My father works for the CIA’). In all these cases, accessing and spreading falsehoods does not have the downstream consequence of leading to true information being divulged, hence no privacy losses.

⁵⁵Ruth Gavison, ‘Privacy and the Limits of Law’ (1980) 89 *The Yale Law Journal* 421, 431.

⁵⁶Ruth Gavison, ‘Privacy and the Limits of Law’ (1980) 89 *The Yale Law Journal* 421, 432. Footnote. Emphasis added.

⁵⁷*ibid.* I will return to Gavison’s notion of attentional privacy at more length in the following section.

⁵⁸*ibid.*

⁵⁹Véliz (n 2) 218.

⁶⁰*ibid.* Véliz provides the example of Iris, who divulges she is celibate after being faced with false claims about her sexual activities.

⁶¹*ibid.* 219.

⁶²*ibid.*

Finally, if John did feel compelled to defend himself against false beliefs about his father's criminal status by divulging truths, it seems he might just do so by parsimoniously stating 'the claim that my father is a convicted murderer is false'. Does such a limited correction diminish his privacy? One might take the view that facts about one's parent's criminal status is ipso facto personal information, as a matter of genre. So the fact that John's father is *not* a convicted murderer is equally a matter of privacy. But this view misses that it is often a proposition's particular content – rather than its mere subject matter – which grounds our interest in keeping relevant information unknown.

We tend to hold a range of defeasible presumptions – quietist default beliefs about others – that we take to be 'normal' or 'neutral'. These presumptions often lack statistical justification; they are not necessarily based on fact. For example, as a matter of default, I tend to presume that any one of my colleagues does not suffer from cancer, that they are not presently depressed, that they are not victims of domestic violence, and that their parents are not convicted murderers. Given that, statistically, a relevant number of my colleagues are bound to have cancer, be depressed, etc., my quietist default beliefs are epistemically unjustified.⁶³ We tend to make these presumptions about others as a matter of social convention. It would seem impolite or intrusive to presume that some of my female colleagues are victims of domestic violence, even though I am more likely to hold correct beliefs by guessing so than by ignoring the possibility⁶⁴. There is, in other words, a socially mandated benefit-of-the-doubt bias in favour of 'neutrality' that governs what we believe about others. These conventional 'neutral' presumptions are part of our public personas: I hold these default beliefs about others and assume that others hold them about me. When information is revealed which merely validates these quietist default beliefs – information that confirms I really am, in this sense, *normal*, such as the fact that my father is *not* a convicted murderer – then it is hard to see how privacy has been lost. (Or, at the very least, it is hard to see that privacy has been lost to a comparable degree as if it became public that my father was indeed a convicted murderer. In other words, facts about by father's criminal status weigh on my privacy in radically different degrees, depending on their content.) If it was publicly revealed that my father is *not* a convicted murderer, I would not take my privacy to be diminished, as this information would only confirm what others already tacitly took to be true about me. (I expect people would respond by saying something like: 'Oh, I did not think he was a murderer anyways'.) So, in correcting other people's false beliefs, it does not seem clear that John would diminish his privacy in a relevant sense.

The privacy losses pointed to by Gavison and Véliz are, at most, indirect,⁶⁵ so the insight that accessing and spreading falsehoods may lead to privacy losses down the line does not explain the intuitive wrongfulness of the actions in *Murderer 2* and *3* in hypothetical variants where no sensitive truths end up being divulged.

6. Do falsehoods diminish 'propositional' privacy or 'attentional' privacy?

Pierre Le Morvan is very motivated to hold that falsehoods can diminish privacy (he gives us eight arguments that suggest excluding falsehoods from information protected

⁶³Conservative estimates suggest that statistically, one in three women and one in ten men has experienced domestic violence <<https://www.ncbi.nlm.nih.gov/books/NBK499891/>>. Nearly a third of adults in the U.S. report having been depressed in their lifetimes <<https://news.gallup.com/poll/505745/depression-rates-reach-new-highs.aspx>>.

⁶⁴I am not here suggesting that the social norms that govern the presumptions we make about others are morally valuable. I simply contend they exist and that we form privacy-related expectations in their light.

⁶⁵Blaauw (n 1) 176.

by privacy is counter-intuitive). In order to offer an account of privacy that is properly viewed as ‘diminished’ when falsehoods are accessed or disseminated, Le Morvan proposes to distinguish between *propositional privacy* (i.e. knowing, or ability to know, of a proposition) and *factive privacy* (knowing, or ability to know, that a proposition is true)⁶⁶. He illustrates the distinction in the following way:

‘Suppose that p is the personal proposition that S carries the BRCA1 and BRCA2 genes predisposing S to a higher than normal likelihood of developing breast cancer. S has propositional privacy relative to me and to p as a function of my ignorance (or inability to know) of p, whereas S has factive privacy relative to me and to p as a function of my ignorance (or inability to know) that p is true’.⁶⁷

Falsehoods only diminish propositional privacy, since they do not enable those who believe them to know that the relevant proposition is true.⁶⁸ Applied to the *Murderer* cases, my reading the letter and disseminating that John’s father was a convicted murderer would leave John’s *factive* privacy intact, but it would diminish his *propositional* privacy because it would entail that I and others form propositional attitudes about John’s father’s criminal status. Does Le Morvan’s account help us explain in virtue of what my actions in *Murderer 2* and *3* were wrongful?

The point of propositional privacy seems to be that people are forming thoughts about an issue at all. This echoes Gavison’s account of ‘attentional privacy’ on which ‘[a]n individual always loses privacy when he becomes the subject of attention.’⁶⁹ To be sure, Gavison suggests – sensibly, in my view – that turning our attention to others by concentrating our thoughts on them (i.e. ‘discussing, imagining or thinking about another person’⁷⁰) is ‘related to privacy in a more indirect way, if at all . . . For the most part, however, thinking about another person, even in the most intense way, will involve no loss of privacy to the subject of this mental activity’.⁷¹

By contrast, according to Le Morvan, propositional privacy is diminished by ‘any propositional attitudes (including speculation)’⁷². Do we have a right that people not entertain any propositions about sensitive facts about us? This seems like an exceedingly broad definition of privacy; on its view, I already diminished my colleagues’ privacy a few paragraphs ago when I wondered whether some of them might be suffering from cancer, depression, or other ills.

In defining propositional privacy such that the mere entertaining of a proposition constitutes a privacy diminishment, Le Morvan effectively equivocates my reading of John’s mother’s letter and then telling others about it (which seems intuitively wrong) and my mere pondering whether or not my colleagues are healthy (which seems intuitively fair game). In other words, yes, propositional privacy counts accessing and spreading of falsehoods as diminishments of privacy. But the stakes of Le Morvan’s account are so deflated that he is unable to explain that diminishing privacy is wrong at all. In fairness, Le Morvan’s objective is to provide a normatively neutral account of

⁶⁶Le Morvan (n 6) 321.

⁶⁷*ibid* 316.

⁶⁸It might even be argued that falsehoods enhance or preserve factive privacy because, unless they evoke suspicion, they lead people *further away* from the truth. Gavison makes a similar point: Gavison (n 55) 421.

⁶⁹*ibid* 432.

⁷⁰*ibid*.

⁷¹*ibid* 432, 433.

⁷²Le Morvan (n 6) 322.

privacy⁷³; he explicitly states that he does not wish to ‘engage in arguments concerning what is, or ought to be, moral and/or legal’.⁷⁴

Propositional privacy is too broad and overinclusive to work as a normative account of privacy/an account of a right to privacy. It therefore is unable to explain what, if anything, was wrongfully committed in the cases of *Murderer 2* and 3.

In the following two sections, I offer two alternatives for explaining the wrongfulness of accessing and disseminating falsehoods. These are choate offences (section 7) and weakening the truth requirement (section 8).

7. Letting attempts account for wrongfulness?

I suspect we are so motivated to find wrongdoing in the *Murderer* cases and in Kappel’s *Adam* in part because the relevant offenders clearly act with malicious intent.⁷⁵ But this does not necessarily mean privacy was diminished with view to the propositions that Adam is a member of the Neo-Nazi Party and that John’s father is a convicted murderer. We could bite the bullet and view that in the cases of *Adam* and *Murderer 2* and 3, no privacy diminishment occurred (and that thus, no privacy rights were violated), but still hold the actions of reading the letters (and telling others about it) constitute wrongs because they are *attempted* privacy violations. We could, in other words, let inchoate offences do the work of capturing the intuitive wrongness of the cases of accessing and spreading falsehoods.

Inchoate liability is a primarily a tool within (the philosophy of) criminal law, so one might wonder whether it is broadly applicable to moral wrongs⁷⁶. However, arguments that morally justify imposing liability for inchoate offences – i.e. the view that inchoate offenders express especially morally blameworthy mental states⁷⁷, and/or because that they create unreasonable risks⁷⁸ – clearly apply to wrongdoing in general. In the *Murderer* and the *Adam* cases, the relevant intruders clearly fit the paradigm for inchoate offences; they clearly intended to violate the victim’s privacy and furthermore did all they could to bring about the relevant harm, so the fact that their actions did not result in choate privacy invasions was simply a matter of *luck*. In the words of Thomas

⁷³One problem of philosophical accounts, which purportedly deal with ‘neutral’ privacy diminishments (i.e. which profess to remain silent about when such diminishments count as privacy *violations*), is that they still get at privacy by asking – normatively – when we take it to be infringed by others. For example, many of the examples of ‘privacy losses’ offered by Le Morvan seem to us such compelling cases of ‘losses’ because they clearly constitute *violations* (opening of letters and diaries, a doctor divulging medical information; *ibid* 319.) By contrast, cases where we can be sure there was no *violation* because access was obviously permitted or excused (i.e. when information is divulged voluntarily), tend to raise the question of whether privacy was ‘lost’ after all. (For a critical discussion on the “voluntary divulgence objection” see Menges (n 53) 388.)

I suspect that supposedly ‘neutral’ accounts of privacy are inescapably normative at their core. What distinguishes generic information about us from ‘private’ or ‘personal’ information, if not the (normative) notion that ‘private’ information would us more socially vulnerable if others were to come by it, or that we have a (normatively) legitimate interest in concealing it? Unless we maintain that “privacy is simply what people say it is”, we have no way of getting around normative conceptions of privacy.

⁷⁴Le Morvan (n 6) 330.

⁷⁵The way Kappel sketches his case of *Adam*, it also seems likely that the intruder committed various other wrongs (the stealing, the breaking in) in the course of their actions. This might partially explain our aversion.

⁷⁶The discussion of whether the behaviour described in *Murderer 2* and 3 is, or ought to be, criminal is beyond the scope of this article.

⁷⁷See Gideon Yaffe, *Attempts in the Philosophy of Action and the Criminal Law* (Oxford University Press 2010) 7.

⁷⁸RA Duff, ‘Criminal Attempts’ 134.

Nagel, it seems intuitive that people should not be morally assessed – positively or negatively – for ‘what is due to factors beyond their control’⁷⁹. The concept of inchoate offences explains moral blameworthiness with a view to situations where such facts which ought not be counted in the defendant’s favour. The chance of the information in the letter being false, or written in a foreign language, should not extinguish moral culpability⁸⁰.

One benefit of using the concept of inchoate offences is that we can derivatively explain the wrongfulness of *Murderer 2* and *3* in reference to characteristics of the full offence of privacy violation. This is intuitively apt; the wrongs committed in the *Murderer* cases (accessing and disseminating falsehoods) seem to have more to do with privacy violations than with some other offence, like defamation. It seems fitting to view those who intrude into falsehoods blameworthy in ways that depend on the wrong of privacy violations. ‘Attempted privacy violations’ conveniently allow us to take the view that some ‘privacy wrong’⁸¹ has occurred, without viewing that privacy was diminished. So far so good: the reading of the letter in *Murderer 2* and the dissemination of information about John’s father in *Murderer 3* – as well as other intrusions resulting in false beliefs – are wrongful because they constitute *attempted* privacy invasions.

However, there are reasons to remain dissatisfied with the ‘attempts’ solution. In particular, the more a case of accessing falsehoods resembles cases of accessing true information, the more it seems arbitrary to let sheer falsehood be the distinguishing factor between ‘merely attempted’ and ‘choate’ privacy violations. Consider a letter from my healthcare provider which is addressed to me and contains a positive Covid_19 diagnosis. My employer opens and reads it without my permission. Based on having read the content of this letter, my employer now forms the belief that I have Covid_19. This is clearly a paradigmatic privacy violation. Still, medical diagnostics remains imperfect⁸². Let us change the hypothetical such that the diagnosis was a false positive⁸³. On this version, I do not, in fact, suffer from Covid_19. Those who hold that privacy cannot be diminished by false beliefs could thus only conclude that my employer *attempted* to invade my privacy by reading my letter. (Since the information about my diagnosis was incorrect, reading it was bound to remain an ‘impossible attempt’⁸⁴.) Does this mean that if someone hacks into a medical database and reads the health information on hundreds of patients, they only violate the privacy of those whose diagnoses happen to be medically accurate (and cases where medical information was incorrect must count as mere attempts)? It seems odd that whether or not my employer committed a choate violation should turn on diagnostic accuracy. This is especially true for cases where the ‘true’ information is never revealed, and everyone, including the ‘victim’, lives on assuming that the diagnosis was correct. In such cases, the falsehood of

⁷⁹Thomas Nagel, ‘Moral Luck’, *Mortal Questions* (Cambridge University Press 1979) 203.

⁸⁰There is a different question of whether there ought to be any difference in culpability, or blameworthiness, between attempts and choate offences. As Grant Lamond notes, whether outcome luck is at all relevant to the defendant’s blameworthiness “has generated the sharpest debates in criminal law theory” Grant Lamond, ‘Criminal Culpability and Moral Luck’ (2021) 23 *Jerusalem Review of Legal Studies* 149, 150. For my purposes here, it suffices that inchoate offences explain that the actions in *Murderer 2* and *3* were at least *somewhat* wrongful (even if they were less wrongful than choate privacy invasions).

⁸¹Compare my discussion of Kappel and Véliz’ accounts in section 23.

⁸²Le Morvan uses a similar point to argue that we ought to view that falsehoods do diminish privacy in some cases, Le Morvan (n 6) 319.

⁸³As opposed to false negative Covid_19 test results, which are more common, false-positive results are extremely rare (but they exist, see Brendan Healy, Azizah Khan, et.al. ‘The impact of false positive COVID-19 results in an area of low prevalence’ *Clinical Medicine*, 2021 Vol 21, No 1: e54-e56).

⁸⁴For a critique, see Yaffe (n 77) 129.

the diagnosis fails to acquire any sociological or psychological relevance. The point here is not that holding those who access falsehoods for attempted privacy violations is misguided in principle, but rather that it does not go far enough: some cases of accessing falsehoods seem to warrant treatment as fully-fledged privacy violations, rather than as inchoate ones.

In consolation to this objection, we might grant that *some* privacy violations are surely choate in the cases of the false diagnosis, because, after all, my employer ended up with *some* true beliefs. Specifically, my employer now knows that (a) I underwent a Covid_19 test and I received a diagnosis from my medical provider (whatever the result of the diagnosis may have been), and (b) that the letter informs me the result was positive (whether or not this is medically accurate). Beliefs with this propositional content remain true. The mistake merely lies at the level of the inference, based on the contents of the letter, that my actual health status was Covid_19-positive.

Of course, these bits of information (a and b) are not quite as sensational as the proposition that I *actually have* a relevant illness. But (a) and (b) are still true beliefs about me that involve sensitive information, so they suffice for a privacy diminishment on my account. Similarly, we might hold that in case of Murderer 2 and 3, I succeeded in diminishing (or violating) John's privacy in virtue of my now holding the (true) belief that *John's mother wrote to him that his father was a convicted murderer* (whether or not this is true). We could piece together the intuitive wrongfulness of my employer reading my medical information – as well as *the Murderer 2 and 3* cases and various other cases of false beliefs – with some patchwork of inchoate privacy violations and choate privacy violations regarding of the less sensational (but true) beliefs. For those who are steadfastly committed to the view that privacy is wholly factive, this mix of inchoate privacy violations and (lesser) choate privacy violations may yield a satisfactory solution.

8. Weakening the truth requirement?

I, however, think the view that accessing falsehoods is merely an attempted privacy violation is unsatisfactory. Intuitively, the choate violations I picked out above do not quite capture the propositional content that is really the subject of privacy-related concerns (i.e. that I *have* Covid_19, that John's father *is* a convicted murderer, etc.). So they seem like a band-aid solution to my normative woes. True, attempts manage to explain the wrongfulness of accessing and publishing falsehoods while tracking the presumption that privacy is factive. But they do nothing to *justify* this presumption: why should the falseness of a piece of information turn a choate privacy invasion into a mere attempt? My example of the false medical diagnosis in the previous section has put pressure on this point. As I argued above, to presume that the truth of the relevant information must make some normative difference (even the mere difference between inchoate and choate versions of the same offence) is to beg the question.⁸⁵

To be sure, it would seem odd to view that any blatant fabrications or rumours would violate privacy,⁸⁶ especially if these are made up out of thin air, without any relevant epistemic work (the digging or searching I mentioned above). I suggested earlier that such blatant falsehoods fit better into the wrong of defamation. Intuitively, when it comes to privacy, there has to be some concern with accessing or propagating things that are true about another person⁸⁷.

⁸⁵See my discussion in section 4.

⁸⁶See Blaauw (n 1) 169. Kappel (n 3) 190. Véliz (n 2) 218.).

⁸⁷Compare also Matheson (n 4) 264.

But it may be sufficient to require that beliefs that are capable of diminishing privacy are *pragmatically* or *conventionally* true; or perhaps just that they are *warranted*⁸⁸ or *rational*. The requirement that beliefs can only diminish privacy if they are *metaphysically true* seems unnecessarily demanding. After all, when it comes to privacy invasions, it is not merely the content of the belief but also *how*⁸⁹ one came to it which makes the normative difference.⁹⁰ Beliefs that come about as a result of random speculation and fabrication typically lack reasons for belief.⁹¹ Merely requiring that relevant false beliefs are warranted or rational – or, perhaps, that they are based on some authentic source – would not lead to an overly broad account of privacy diminishments, because this view only picks out cases where people had some relevant justification for forming their false belief. In these cases, people were at least somewhat concerned with, or effort put towards, getting at truth.

As Kappel and Le Morvan have argued, false information – at least in cases where believing it is justified – can make us equally socially vulnerable as true beliefs.⁹² The moral prohibition against discrimination is a useful analogy: it extends to actual and *perceived* group membership⁹³, so if a person is treated disadvantageously on the assumption that they belong to a protected religious minority, we count this as discrimination (rather than as merely attempted discrimination), even when the assumption about their relevant group membership turns out to be wrong⁹⁴. Including cases of perceived group membership within our understanding of discrimination tracks how discriminatory treatment victimises people.

Weakening the truth requirement has benefits with view to privacy diminishments that cannot neatly be represented in terms of propositional content, like sensory data⁹⁵, accessing evaluative statements like ‘X is dumb and irresponsible’,⁹⁶ or changes in doxastic states, such as suspending a belief. Consider the case of my romantic partner who believes that I love him. After reading my diary without permission, my partner becomes aware of the ambiguous feelings I harbour towards him. He thus tells me: ‘I don’t even know if you love me anymore’. He has here suspended his belief that I love him. Such a suspension of belief is not really propositionally true or false; at most, we can criticise someone who suspends their belief for ignoring the epistemic reasons they *have* for keeping a relevant belief. A wholly factive notion of privacy runs into difficulties in these domains (for this reason, scholars sometimes distinguish between different modes of losing privacy.⁹⁷ Such distinctions complicate philosophical accounts of privacy. Weakening the truth condition may simplify things.)

⁸⁸Munch and Mainz argue that privacy can be diminished by ‘warranted beliefs’. Their account remains intentionally neutral on which kinds of beliefs diminish privacy, and whether such beliefs must be true: Munch and Mainz (n 33) 274.

⁸⁹This point is made by Andrei Marmor, ‘What Is the Right to Privacy?’ (2015) 43 *Philosophy & Public Affairs* 3, 4. and Thompson (n 23) 307.

⁹⁰After all, sensitive facts about others can be known innocently, such as in cases where I accurately infer another person has cancer based on seeing them receive treatment in on the oncology ward of a hospital.

⁹¹It helps that falsehoods typically lack justification. (For a detailed discussion of how much justification is required, and if ‘lucky’ true beliefs could violate privacy, see Véliz (n 2) 227).

⁹²Kappel (n 3) 190. Le Morvan (n 6) 320.

⁹³Tarunabh Khaitan, ‘A Theory of Discrimination Law’ 145.

⁹⁴I discuss analogies between privacy invasions and discriminatory treatment in more detail in [anon].

⁹⁵Véliz (n 2) 235.

⁹⁶Gavison suggests that such statements leads to some loss of privacy, as “information what the speaker thinks about X is also information about X”, see Gavison (n 55) 432.

⁹⁷See, e.g. Véliz (n 2) 235.

Given all these virtues, why do some philosophers like Blaauw take it to be so incontestable that privacy is factive?⁹⁸ When we look at parallel facticity debates – such as debates over whether expressions like ‘information’ is factive⁹⁹, or whether verbs like ‘to remember’, ‘to learn’, ‘to reveal’, and even ‘to know’¹⁰⁰ are factive – we run into a similar ‘intuition stalemate’¹⁰¹. Allan Hazlett has suggested that divergences in intuitions on whether it is possible ‘to know’ a *falsehood* might be down to methodological differences within conceptual analysis. He argues that so-called ‘factive’ verbs, including ‘to know’, are not semantically factive (i.e. they can be used in non-factive ways in ordinary talk)¹⁰². But this, on Hazlett’s view, ought not worry those working in the ‘post-Gettier’ epistemology tradition; for they ‘are concerned with identifying necessary and sufficient conditions for knowledge attribution’¹⁰³, whereas those who employ a ‘linguistic method’ – like Hazlett himself – elicit intuitions ‘concerning whether or not a character in the story *said something acceptable*’¹⁰⁴. These approaches – the ‘traditional epistemological’ one; the ‘linguistic/ semantic’ one; and others, like value-oriented inquiry¹⁰⁵ – represent distinct projects within conceptual analysis.

A similar point can be made regarding the philosophical debate over whether ‘information’ is factive. Proponents of the veridicality thesis (VT) maintain that information must be truthful (‘p counts as information only if p is true’¹⁰⁶). Note that Floridi’s influential veridical conception of semantic information was formulated, at least in part, to address a particular problem that arises in relation to the objective of quantifying the informational content of sentences in formal languages (the Bar-Hillel Carnap paradox¹⁰⁷). VT is confined to the sub-category of *declarative semantic information*. It does not extend to other kinds of semantic content like stipulations, instructions, or invitations. These – as defenders of VT grant – ‘may be correctly qualified as kinds of information’¹⁰⁸, even if they lack truth/ truthfulness. Developing a concept of ‘information’ for these purposes is clearly distinct from the linguistic project of making sense of the term ‘information’ in ordinary language.

Linguistically, of course, term ‘information’ can clearly be used in a more generic, synecdochic¹⁰⁹ sense. In ordinary language, we habitually allow that information can be ‘false’ without insisting that this undermines its ontological status of being ‘information’: anyone who has lived through the Covid_19 pandemic will agree that the expression

⁹⁸I thank my reviewer for prompting investigation into the roots of divergent intuitions on whether privacy is factive.

⁹⁹I mentioned the debate over whether information is veridical in the introduction.

¹⁰⁰Allan Hazlett, ‘The Myth of Factive Verbs’ (2010) 80 *Philosophy and Phenomenological Research* 497, 501.

¹⁰¹John Turri, ‘Mythology of the Factive’ (2011) 2 *Logos & Episteme* 141, 145.

¹⁰²Hazlett argues that when it comes to ordinary talk about knowledge, factivity is merely pragmatically implied, but cancellable Hazlett (n 100) 513.

¹⁰³*ibid* 497, 498. (Emphasis as in original).

¹⁰⁴*ibid*. (Emphasis as in original).

¹⁰⁵Hazlett cites, amongst others, Sally Hallsinger, ‘Gender and Race: (What) Are They? (What) Do We Want Them To Be?’ (2000) 34 *Nous* 31, 33. I discuss value-laden approaches to conceptual analysis further below.

¹⁰⁶Floridi (n 8) 31.

¹⁰⁷Yehoshua Bar-Hillel and Rudolf Carnap, ‘Semantic Inormation’ (1953) 4 *The British Journal for the Philosophy of Science* 147. Details are beyond the scope of this article. I have no view on whether a veridical conception of information deals better with the BCP and related paradoxes than alethically neutral conceptions of information.

¹⁰⁸Floridi (n 17) 83.

¹⁰⁹*ibid* 104.

‘spreading false information’ is coherent and meaningful.¹¹⁰ ‘False information’ is not the mere absence of information (i.e. non-information), but it contains relevant semantic content (which, as proponents of VT grant, might even be somewhat informative¹¹¹). As pointed out by Lundgren, ‘our judgment of whether something is information comes before, or at least is separate from, the judgment whether that something is true (truthful) or false’¹¹². At least in ordinary language, it is permissible to conceive of ‘information’ as ‘alethically neutral’¹¹³. This illustrates that (some) disagreement over the factivity/veridicality of information may be a matter of conceptual analysis method.

To be sure, Hazlett’s claims about non-factive uses of verbs like ‘to know’ has faced challenge even within the ‘linguistic method’ he takes himself to be working in¹¹⁴. But this does not undermine his more fundamental point that differences in intuitions regarding the question of whether a concept is factive may ultimately depend on what exactly it is that analytical philosophers are doing when they engage in *analysis*. When I endorse the argument, as I did in the present section, that we should adopt a non-factive conception of privacy because we need privacy to protect us from the kinds of false beliefs that make us just as socially vulnerable as true beliefs¹¹⁵, I am pursuing what Hazlett calls a value-oriented inquiry and what others have termed ‘conceptual engineering’¹¹⁶ or ‘conceptual ethics’¹¹⁷. I am, in other words, asking: *What do we have the concept of privacy for?* My implicit answer is that privacy is supposed to provide a normative bulwark against the unwelcome intrusions of others¹¹⁸. This protective force is more capacious if it covers (intrusive) falsehoods. So non-factive privacy is better at tracking how intrusions victimize people. To be sure, in proposing this, I take myself to be doing something more modest and less creative than philosophers like Sally Hallsinger, whose approach to conceptual analysis – also dubbed ‘ameliorative

¹¹⁰See Floridi, who grants that ‘[l]inguistically, the expression ‘false information’ is common and perfectly acceptable’ *ibid* 93. (Though Floridi proceeds to suggest ‘quasi-information’ and ‘misinformation’ are more apt terms; *ibid* 95.).

¹¹¹Floridi permits that falsehoods can still be ‘*informative*’. He explains that ‘a false statement *ss* may be *more informative* than a true statement *s’s*’ [...] By way of example, suppose that you are running a catering contract for an event, and that there will in fact be exactly 200 people in attendance. Suppose that *ss* is *there will be 201 people in attendance*, and *s’s* is *there will be between 100 and 200 people in attendance*. *s’s* is true whilst *ss* is false, but *ss* is more informative than *s’s*’ on any natural understanding of the concept INFORMATIVE’ see Sequoiah-Grayson, Sebastian and Luciano Floridi, ‘Semantic Conceptions of Information’, *The Stanford Encyclopedia of Philosophy* (Spring 2022 Edition), Edward N. Zalta (ed.), <<https://plato.stanford.edu/archives/spr2022/entries/information-semantic/>>.

¹¹²Lundgren (n 17) 2893. It strikes me that it is perfectly coherent to ask another person the question: “Are you confident this information is true/ truthful?”. Note that it would be paradoxical to speak of “false facts” and tautological to ask the question “Are you confident this fact is true/ truthful?” This suggests ‘fact’ is factive in a sense that ‘information’ is not.

¹¹³*ibid* 2886.

¹¹⁴The point remains a matter of continuing debate. For accounts that explain why Hazlett’s allegedly “non-factive” uses of ‘to know’ are factive after all, see Turri (n 105) and Wesley Buckwalter, ‘Factive Verbs and Protagonist Projection’ (2014) 11 *Episteme* 391. On the other hand, it has been suggested that factive versus non-factive uses of verbs like ‘to know’ varies significantly across languages; see Roberta Colonna Dahlman and Joost van de Weijer, ‘Cognitive Factive Verbs across Languages’ (2022) 90 *Language Sciences* 101458.

¹¹⁵Similar points have been made by Kappel (n 3) 190. Le Morvan (n 6) 320.

¹¹⁶Steffen Koch, Guido Löhr and Mark Pinder, ‘Recent Work in the Theory of Conceptual Engineering’ (2023) 83 *Analysis* 589.

¹¹⁷Alexis Burgess and David Plunkett, ‘Conceptual Ethics I’ (2013) 8 *Philosophy Compass* 1091, 1094.

¹¹⁸That I view privacy as inherently normative should come as no surprise if you have read my footnote 73.

inquiry'¹¹⁹ – proposes to focus our efforts not on capturing what we do mean, but on 'how we might usefully revise what we mean for certain theoretical and political purposes'¹²⁰. Arguably, privacy is already ambiguous and non-factive uses find at least some support in the literature¹²¹, so in making the case that the non-factive version is normatively preferable, I do not quite propose *revisions* to the concept of privacy. Still, I am suggesting that 'moral considerations are amongst those that can be ought to bear'¹²² on which of the available concepts of privacy we should pick.

Going back to Hazlett, I grant that those engaged in 'traditional epistemology' might not care for my normative arguments in favour of non-factive privacy. For those (perhaps Blaauw is amongst them) who take it to be *analytically* true that falsehoods cannot diminish privacy, arguments about the protective force of non-factive privacy will not move the needle. Those 'traditional epistemologists' might be content with the solution I offered in the previous section, i.e. the view that intrusive ways of accessing falsehoods can count as *attempted* privacy violations. But some (including this author) might lack antecedent commitments to the notion that privacy is factive. For us, weakening the truth requirement for privacy diminishment is an attractive alternative.

9. Conclusion

My argument in this article is animated by the objective to explain the wrongfulness I commit when I access false information contained in a private letter, form relevant false beliefs, and then disseminate these (see the cases of *Murderer 2* and *3* in my introduction). Both access and control accounts of privacy (at least insofar they view privacy as factive) would deny that such instances of accessing and spreading false information constitute privacy diminishment. If we presume, as I do, that the right to privacy can only be violated via privacy diminishment, then I have to deny that cases of accessing and disseminating falsehoods are privacy violations. But how else to explain the wrongfulness of these acts?

I have argued that the academic literature has not provided a satisfactory answer to this question. The concept of defamation does not intuitively get at the wrong at hand. Accounts that confirm privacy violations even in cases where privacy remains undiminished sketch the right to privacy too broadly, and the notion that falsehoods may pressure privacy losses down the line is too indirect to explain in virtue of what accessing falsehoods is wrongful. Finally, the notion of 'propositional' privacy is too normatively inert to explain that any such 'propositional' privacy diminishment constitute pro tanto wrongs.

I have offered two alternative ways of explaining the wrongfulness of accessing and disseminating falsehoods: those committed to the notion that privacy is factive may like to think of instances of accessing falsehoods as 'attempted' privacy violations, and those who take a more liberal antecedent view on whether privacy can be diminished by falsehoods might prefer to weaken the requirement that relevant information must be metaphysically true. If we follow Hazlett, we can view these as different solutions for different projects (and methodologies) for conceptual analysis: factive privacy may suit

¹¹⁹Katharine Jenkins, 'Amelioration and Inclusion: Gender Identity and the Concept of Woman*' (2016) 126 *Ethics* 394, 359.

¹²⁰Halslanger (n 105) 34.

¹²¹In defensive of factive privacy, see, e.g. Blaauw (n 1). In defence of non-factive privacy, see, e.g., Le Morvan (n 6).

¹²²Burgess and Plunkett (n 117) 1094.

‘traditional epistemologists’, whereas non-factive privacy seems a more suitable target concept¹²³ for those who seek to use the concept of privacy for normative purposes.

Alice is Lecturer in Law at Stanford Law School. She was previously Assistant Professor of Law at New College of the Humanities, London. Alice holds a DPhil in Law from Oxford University, a B.A. in Politics, Psychology, and Sociology from Cambridge University, and degrees in Law from King’s College London and the Humboldt University in Berlin. <https://orcid.org/0009-0009-6555-2459>

¹²³Jenkins (n 119) 395.

Cite this article: Schneider A. (2025). “Can privacy be diminished by falsehoods?” *Episteme* 1–21. <https://doi.org/10.1017/epi.2025.13>