The Moral Significance of Privacy Dependencies

Authors: Lauritz Aastrup Munch (Aarhus University) and Jakob Thrane Mainz
Contact: (lauritzmunch@gmail.com, jakob-mainz@hotmail.com).
Forthcoming in *Philosophy and Technology* (please cite the published version).
Penultimate version

Abstract

Often, when we share information about ourselves, we contribute to people learning personal things about others. This may happen because what we share about ourselves can be used to infer personal information about others. Such dependencies have become known as 'privacy dependencies' in the literature. It is sometimes claimed that the scope of the right to privacy should be expanded in light of such dependencies. For example, some have argued that inferring information about others can violate their right to privacy. Others have argued that sharing personal information about yourself that license such inferences can by itself violate the right to privacy. In this paper, we argue that the latter view should be rejected.

keywords: privacy rights; privacy dependencies; fairness; waivability; transparency; genetic testing

I. Introduction

Recently, Peter became interested in his family's genealogy after hearing incredible stories about the life of his great grandfather. He therefore decides to send a DNA sample to 23andMe—a company specializing in ancestry research—hoping that they can help him put together a comprehensive overview of his family's history. Peter discovers that 23andMe also specializes in genetic health testing. Having studied the life of his family members several generations back, he knows that heart diseases tend to run in the family. Peter therefore decides to order a genetic risk profile from 23andMe. A few months later he receives the results. With horror, he learns that he has a hereditary predisposition for a certain heart disease. He immediately posts the sad news on his social media profile. Within minutes, Peter receives a phone call from his identical twin

brother Carl. He is furious and accuses Peter of *having violated his right to privacy*. He tells Peter that he did a similar test years ago. But he never told anyone, because doing so would reveal that Peter—being his identical twin—has the same predisposition, and he therefore feels that the information is not just his to share. Peter defends himself, arguing that he has—*also due to having a right to privacy*—a right to share any personal information about himself that he wishes, and that it is just too bad that doing so prevents Carl from concealing that he has the predisposition too.

The structure of this case is such that only one of the twins can have it their way. And since both brothers have strong reasons supporting their favored outcome, we face what seems like a moral dilemma. Both brothers defend their positions based on their right to privacy. Peter claims that his right to privacy includes the right to share personal information about himself with others. This is indeed a commonsensical idea, and one that enjoys much support in the literature.[1] Carl, on the other hand, claims that his right to privacy includes that others do not access personal information about him in ways that he has not authorized.[2] And when Peter shares his genetic information, it is made all too easy for others to access Carl's personal information via simple rules of inference in a way that short circuits his authority over the access to the information. But it is unclear if we should side with Peter or Carl here (or perhaps flip a coin?), at least from the perspective of the right to privacy.

Of course, this paper is not about Peter and Carl in particular. Rather, our concern is how to think about the scope of the right to privacy (and resulting duties) under conditions where what you share about yourself can be used to reliably infer things about others. Such 'inferential' relations are foregrounded in the case above due to the genetic identity of identical twins. This is an example of what has become known as 'privacy dependencies' in the literature (Barocas and Levy 2020). More carefully, then, our question is if the existence of privacy dependencies ever makes a difference to Peter's permissions with regards to what information he

[1] See, e.g., Thomson (1975), Marmor (2015), Hanin (2022), Rumbold & Wilson (2019), Moore (2018)
[2] See Westin (1968), Fried (1968), Moore (2003, 2010), Gross (1971), Parker (1974), Parent (1983), Allen (2003), Rössler (2005), Lundgren (2020), Menges (2021).

may share about himself. Answering this question is both theoretically important and practically urgent, and this is probably easiest to see by pointing out that privacy dependencies are especially pervasive in the age of Big Data. For instance, a class of privacy dependencies have come to our attention thanks to research suggesting that state-of–the-art algorithmic systems relying upon statistical techniques and fueled by Big Data have proven capable of inferring sensitive personal information from mundane sources. One intriguing finding is Wang and Kosinski's (2018) paper showing that it is possible to accurately predict sexual orientation based on facial pictures. Another famous example is Gebru et al.'s (2017) paper showing that it is possible to predict people's political preferences based on the type of car they drive. What this means is that when you share pictures displaying your face online, or pictures of the car you drive, you easily risk contributing to the (training of the) predictive models that make it harder for people who look like you—or drive the same type of car as you do—to conceal their sexual orientation and political preferences. This could be taken to imply—if Carl is right—that we routinely contribute to causing a threat to other people's privacy in ways similar to how Peter might be said to threaten Carl's privacy.[3]

Our aim in this paper, then, is to make progress on the question of how to think of the moral significance of privacy dependencies. But we are not going to tackle the question by arguing for or against the view that Peter wrongs Carl. The reason for this is that the genetic case seems to us a *critical case* for the view that privacy dependencies ever could make a difference to what we are permitted to share. If you do not find it *prima facie* compelling that Peter should desist for the sake of Carl's privacy (or that there at least is a moral dilemma here), you probably don't think that privacy dependencies could ever matter to the scope of the right to privacy.[4] Our strategy will instead be to ask the question of how to move from this localized verdict to a more general view of privacy rights and duties that take into consideration privacy dependencies. And

[3] More technically, the algorithm is discovering and subsequently applying rules that, if generally known, would make it much harder for people to keep things concealed that could be inferred from mundane information. It is the discovery of these rules that people are contributing their information to.

[4] See Véliz (2020: 77).

as we will show below, there are deep structural problems associated with this move that ultimately suggest to us the following: If privacy dependencies matter to the scope of privacy rights, they will only do so in very special cases—the genetic case being one such. Hence, our move is to offer two arguments that block generalizing the moral insight there is to be had from the case with which we began.

II. Precisifying the problem

Our question is if the existence of privacy dependencies can affect what people are permitted to share about themselves due to a concern for protecting the privacy of others. But before we can tackle this question, it is helpful to present in greater detail three claims that engender dilemmas of the sort with which we began:

> *Excludability*
> The right to privacy encompasses A's claim to exclude C from accessing A's personal information p_A.
>
> *Waivability*
> The right to privacy encompasses A's moral power to permit C to access personal information p_A by sharing p_A with C.
>
> *Privacy Dependency*
> When A shares personal information p_A with C, it can make it possible for C to infer (and in this way access) personal information p_B about B.

The two first claims are normative claims about the structure of the right to privacy. Each identifies a so-called 'incident" that should be familiar from the Hohfeldian (1919) analysis of rights, and indicates what it ranges over in the context of the right to privacy. The third claim is a descriptive statement saying that *a* way of accessing personal information is via inferential means.

Excludability takes the form of a Hohfeldian *claim*. If person A has a claim over something, it implies that another person B has a duty of non-interference with regards to this

something.[5] A common example would be the claim somebody has over their rightful property. Because one has this claim, others have a duty not to interfere with, manipulate or otherwise control their property. We call this idea "Excludability" because others are excluded, morally speaking, from engaging in certain ways with the object (or domain) that the claim attaches to. In the context of the right to privacy, the kind of exclusion at stake is typically exclusion from *accessing information in certain ways*. The basic idea that the right to privacy encompasses *Excludability* in this sense is common in the literature. If the right to privacy exists at all, it seems, it must at least encompass the right to exclude others from accessing personal information about oneself.[6]

The second claim—*Waivability*—is modeled over a Hohfeldian moral *power*. Powers are higher-order incidents in the sense that permit the rights-holder to manipulate lower order incidents at will. Common examples of such moral powers are the promissory power (the power to create duties for oneself to fulfill a promise), the power of command (the power to create duties for others) and the power to consent (the power to create a permission).[7]

In the context of the right to privacy, the relevant moral power is the power to *waive* the claim identified in *Excludability*. A waiver—again glossed in Hohfeldian terms—means creating a permission for others to do things (where there would typically otherwise be a duty not to do those things). For instance, if you have a duty not to read my diary—something that seems to follow from *Excludability*—I can exercise my power of waiver and make it permissible for you to do so, for instance by inviting you to do it. *Waivability* is deemed an integral part of the right to privacy by many.[8] It is important to see that waivers can take many forms in practice, and we shall focus only on a subset here for ease of exposition. In the context of the right to privacy, some waivers (of claims to not having one's information accessed) come about through sharing

[5] Hohfeld (1919).
[6] See Rumbold and Wilson (2019) and Mainz and Uhrenfeldt (2021).
[7] Though notice that the normative landscape can be manipulated by other things than exercises of will. For instance, claims can cease to exist due to so-called *forfeitures* (Hanin 2022) and new claims may arise as a result of *wrongings*.
[8] See Thomson (1975) for a canonical treatment, but most seem to agree that the right to privacy comes with the power of waiver (Moore 2018).

the information that one means to waive a claim against having accessed. Let's call this type of waiver "waiver-by-sharing".[9] This is meant as a technical term. If I tell you a secret about myself, I waive my right to privacy with regards to my secret against you. But if I instead invite you to read all the content in my diary, and my secret is written in the diary, then the information conveyed in the speech act that constitutes my waiver ("Please, go on and read my diary") comes apart from the information that I waive my right to privacy over. Call this second type of waiver "waiver-by-declaration". We shall focus mostly on waivers-by-sharing here, but nothing of substance hangs on which type of waiver is involved.

Here are two further points that apply to both *Excludability* and *Waivability*. First, the claims are meant to describe only the structure of the right to privacy, and so do not specify its scope. As a result, they are neutral on whether inferential access to information ever falls within their scope. This is, after all, what we are exploring here.[10] Second, we assumed above that the kind of acts that the right to privacy marks out as morally impermissible (due to *Excludability*) and whose moral status can be manipulated (due to *Waivability*) are acts of accessing personal information.

Nowadays, though, many seem to prefer a more contextual or relational conception of privacy and one might reasonably wonder how our remarks would fit in such a framework.[11] To answer this question, consider for purposes of illustration Helen Nissenbaum's influential reworking of the concept of privacy as 'contextual integrity' (2010). One central idea found in Nissenbaum's work is that in every society we observe a number of social norms that serve to specify what an appropriate flow of information between contexts and between agents looks like. For instance, according to widespread norms governing friendship, it is expected that pairs of friends can share intimate information with each other in the expectation that this information won't be disclosed publicly by the other party (Nissenbaum 2010: 146; Rachels 1975). And when

[9] Hanin (2022) calls such waivers 'implied waivers'.
[10] Scanlon (1975) claims that the scope of the right to privacy has a conventional scope whereas Marmor (2015) suggests that this cannot be the full story.
[11] We thank an anonymous reviewer for prompting us to address this question.

friends share intimate information in accordance with the relevant norm, privacy is not lost or diminished. Rather, to borrow an idea from Julie Inness, we might say in such cases that "we are including another within our realm of privacy, not lessening our privacy" (Inness 1992: 46). Privacy, according to this picture, is only lessened or diminished when information flows in ways that violate relevant information norms and thereby undermine what Nissenbaum calls 'contextual integrity'. The simple framework for thinking about the right to privacy we presented above is compatible with the contextual conception of privacy as it only sketches a bare-bones structure of this right. For instance, the claim being picked out by *Excludability* may be glossed as the idea that one has a claim that one's information is not shared or accessed in ways that contradict social informational norms.[12] *Waivability* also fits within this framework as many informational norms have as part of their content that people should decide for themselves with whom their information is shared.[13] This can also be illustrated with the friendship example presented above: There wouldn't be anything problematic about you (in contrast to your friend) disclosing your information on social media and, plausibly, the friend would be permitted to share your information on social media if you gave them permission first.[14]

Finally, let's consider the third proposition, namely *Privacy Dependency*. This is a non-normative proposition that states, roughly, that if I share a piece of personal information with a third party, it might be possible for that third party to infer a piece of personal information about you. The empirical truth of *Privacy Dependency* has been established many times in the literature.[15] But it's important to note that privacy dependencies can come in many shapes. This is partly because inferential rules come in different shapes and partially because of the different ways in which information about some individuals—as a matter of brute fact—are

[12] This strikes us as coming very close to Scanlon's (1975) conception of the right to privacy.

[13] See Nissenbaum 2010: 148.

[14] Some claim that there exists such a thing as *group privacy* (e.g., Floridi et al. 2017). According to this picture, groups can have emergent forms of privacy and even emergent privacy rights that are not simply the privacy (and privacy rights) of the individuals composing the group. Such a conception of privacy rights can also be made sense by using the framework we have introduced above, since the Hohfeldian framework merely describes the structure of rights and doesn't by itself tell us what rights that exist.

[15] Barocas & Levy (2020), Véliz (2020), Mühlhoff (2021), MacCarthy (2011), de Brouwer (2020), Fairfeld & Engel (2015).

'fertile', that is, apt for being used as premises in inferences that reveal something about other individuals. For ease of exposition, we can distinguish two common types of inferences: truth-preserving inferences and non-truth preserving inferences. Here's an example of a valid truth-preserving inference that we are already familiar with:

Genetic Predisposition

(α) Peter has a genetic disposition for heart disease X and Y.

(β) All identical twins have the same genetic dispositions.

(γ) Peter's identical twin Carl has a genetic disposition for heart disease X and Y.

The inference from α and β to γ is truth-preserving, because the truth of α and β guarantees the truth of γ. It's impossible for α and β to be true while γ is false. And this is the case regardless of whether α and β are actually true. Therefore, the information that "Peter's identical twin Carl has a genetic disposition for heart disease X and Y" is in some sense already contained in the combination of information that "Peter has a genetic disposition for heart disease X and Y" and "all identical twins have the same genetic dispositions."

This is not the case for non-truth preserving inferences like the following, which we are also already familiar with:

Political Preference

(δ) Peter drives a car of type Z.

(ε) There's a .9 correlation between driving a car of type Z and voting Republican.

(ζ) Peter votes Republican.[16]

[16] This example is inspired by an example from Mainz (2022).

Here, the truth of δ and ἐ don't guarantee the truth of ζ. Just because driving a car of type Z is highly correlated with voting Republican, it does not follow that if Peter drives a car of type Z, he must vote Republican. Of course, it may be statistically likely that he votes Republican given his choice of car, but he might as well be among the 10% who own the same type of car but don't vote Republican.

It's important to note that when one waives one's right to privacy with regards to certain pieces of personal information, by sharing the information with others, one can contribute to other's lack of *Excludability* in certain ways. In Genetic Predisposition, Peter contributes the crucial piece of information that he has a genetic predisposition for heart disease X and Y. This is a crucial piece of information because it delivers one of the premises that makes it possible to infer, against his will, that Peter's identical twin Carl has a genetic predisposition for heart disease X and Y. But things are different in Political Preference. Here, Peter also delivers one of the core premises that makes the inference possible. But—assuming Peter actually votes Republican—he also contributes to establishing the correlation between driving a car of type Z and voting Republican. When people worry about privacy dependencies in the era of predictive analytics and Big Data, it's often because of the incredible inferential power these technologies hold. Even though the inferences generated by these technologies are mostly non-truth preserving (like the one in Political Preference), they are still cause for concern. Partly because the statistical predictions generated by these technologies can be impressively accurate, and partly because they can find statistical patterns that were otherwise invisible to the human eye. One thing is to be worried that your identical twin shares their DNA information with others, who then make simple deductive inferences about you in their mind. But another thing is to be worried about how personal information can be inferred from personal information about yourself that seems completely unrelated to the inferred information. Below, we'll address how to think about such privacy dependencies.

III. Against the moral significance of privacy dependencies

With a better grasp of the underlying propositions that generate the kind of dilemmas that interest us, we can now proceed to the substantive moral question of how such dilemmas could be resolved. But since our arguments below will concern the question of what determines the appropriate scope of a right, we must assume some substantive views on what fixes the scope of rights in general (for discussion of this see for instance Raz 1986). There is one necessary condition for the justification of a right that will play a central role in our argument, so it is worth presenting it in detail. This is the idea that a right must be scoped in a way such that the facts that shape people's resulting duties are facts that are generally accessible to them. Following Bolinger, we can refer to this idea as:

> *Transparency*
>
> The scope of a right is justified only if the facts that shape people's corresponding duties are facts that are in principle accessible to them (see Bolinger 2021: 6).

Bolinger (2019; 2021) claims that we should endorse *Transparency* because scoping rights in ways that do not respect it invariably leads to an unfair distribution of the *costs of errors* when agents acting in good faith seek to comply with their moral duties. The relevant errors, and their corresponding costs, come in two types: *false-positive errors* and *false-negative errors*. From the perspective of duty-bearers, a false-positive error occurs in cases where the duty-bearer acts as if (and having the belief) that they have a permission but do in fact *not* have such a permission. As an illustration, imagine that Al gets the impression that Ben consented to him borrowing his car, but Ben didn't in fact. Al now suffers the cost resulting from his false-positive error of becoming a wrongdoer and therefore becoming liable for rectificatory purposes. On the other hand, *false-negative errors* occur when a duty-bearer acts as if they lack a permission to do something (and

believe they lack such a permission) but in fact has the permission in question. In such cases, the cost of error falls on the duty-bearer in the form of the forgone benefit or opportunity cost. Imagine, differently, that Al got the impression that Ben wouldn't want him to borrow his car, but that Ben in fact permitted this (but perhaps forgot communicating it). In this case, Al suffers the opportunity cost associated with a *false-negative error*. We can imagine structurally similar errors and corresponding costs from the perspective of the rights-bearers happening via the exercise of a moral power. In such cases, we shouldn't measure the errors by considering if an action or inaction was permissible or impermissible, but instead by looking at whether the result of the exercise of a moral power corresponds to its intended aim. Take consent as an example. A *false-positive error* would occur in cases where a consentee issued a genuine consent without intending this.[17] (Imagine that somebody uttered some words in a foreign language, taking them to mean "I'd like a cup of coffee" but in fact it meant "I'd like you to paint my car"). The cost would be permitting something that the consentee wouldn't want to permit upon reflection. A *false-negative error*, on the other hand, could come about were a consentee to omit a speech act because they suspected it would token consent, where it in fact wouldn't have. The cost here would be omitting incidentally valuable acts needlessly. (Suppose A avoided declaring "I Love You" because they falsely believed this would amount to issuing a consent to sexual intercourse).

It's probably impossible to fully eliminate the risk of false-positive and false-negative errors when right-holders and duty-bearers seek to coordinate their conduct, even if all parties are acting in good faith. But it is a problem if duty-bearers are systematically deprived of access to information necessary to reliably avoid making these errors. In such cases, they will be imposed costs due to making errors they had little chance of avoiding and such costs are generally unfair. This will happen if the facts that determine the relevant moral statuses are typically inaccessible to them. In such cases, duty-bearers *"face a high stakes gamble with non-trivial risks of error"* (Bolinger 2021: 5). If it turns out that a given conception of some right implies this,

[17] Not all views on consent allow for this. See Bolinger (2019) for discussion.

we have a good reason to reject its specification. The two following subsections will consider the normative case for scoping *Excludability* and *Waivability* in light of *Privacy Dependencies*, respectively. We hope to show that neither should be scoped to take into consideration *Privacy Dependencies* in a way that affects whether people are permitted to share information about themselves.

III.a. Excludability

Recall the example with which we began, the case of the two identical twins and the conflict over whether genetic information should be shared or not. We are now in a position to better appreciate that there is a moral conflict in this case. We must grant that *Excludability* is to be interpreted broadly to include an entitlement to prevent access to information via inferences. If one were to deny this claim, there would be no basis for the complaint that Peter violates Carl's right to privacy by disclosing their genetic information. This is because Peter's genetic information by itself doesn't reveal anything about Carl's genetic information. Carl is only affected when Peter's genetic profile is combined with the information that Carl is Peter's identical twin, and that identical twins have identical genetic profiles. So, no tension would arise in the first place. Now, consider the following argument:

The Narrow Excludability Argument

Premise One: A plausible specification of a Hohfeldian claim requires that the costs of errors are justifiable to duty-bearers.

Premise Two: Interpreting *Excludability* to include restrictions on inferential access to information means that duty-bearers will be liable to make errors that cannot be justified to them.

Conclusion: Interpreting *Excludability* so as to include restrictions on inferential access to information amounts to an implausible specification of a Hohfeldian claim.

Premise One relies on the idea introduced in the former section that costs of errors must be justifiable to duty-bearers. We have seen why something in its vicinity is plausible, so let's focus next on Premise Two. It suggests that interpreting *Excludability* to include restrictions on inferential access puts duty-bearers in a position where they will be imposed unacceptable error-related costs.

Let's argue in favor of this point. When we say that duty-bearers must take into consideration not only what they are sharing with others, but also what others could become capable of inferring about others in part based on what we share, we are demanding that duty-bearers must take something into account they mostly cannot. This is because what a recipient can infer from information I disclose to them depends in part on their cognitive capabilities and, more importantly, what additional pieces of information the recipient has available to them. We simply cannot be aware of these things in most cases and this makes putative duty-bearers prone to suffer costs of errors.

Consider an example to appreciate this point. Suppose Miles is concerned for his brother and chats to a friend about their concern (the brother has behaved weirdly as of late) and that Miles' friend happens to have recently read a lot of medical journal papers. Miles is careful not to reveal health-related information—or any other privacy-sensitive information—about their brother. However, it just happens to be the case that the friend can infer a lot about Miles' brother's health based on how Miles describes the thing that concerns them. This would be a case in which—if we say that *Excludability* should prohibit sharings of information that license certain inferences—Miles wronged his brother and would therefore be liable to bear rectificatory costs. However, Miles had little reason to think he would wrong his brother and we can therefore classify the case as a *false-positive error*. Miles acted against their duty, thinking he acted permissibly.

We'd like to make two points based on this case. First, we don't think it's fair to count Miles as a wrongdoer because he shouldn't bear the costs for making an error he couldn't reasonably avoid. Secondly, we think this point generalizes because the uncertainty that Miles faces is not a contingent fact about the case; rather it is due to the nature of the facts that it would seem that Miles would have to take into account (this is why Transparency, speaking of principled uncertainty, is relevant here). Most of the time, we won't know in great detail what additional information recipients have available, and we won't know in great detail what inferences their cognitive capacities allow. The relevant facts are internal to the minds of the, and since we aren't reliably capable of mind-reading (most of the time we even lack reliable introspection into our own cognitive abilities and beliefs), we are not reliably capable of acting in a way to robust avoid sharing information that will be used to infer things about others. This suggests to us that, generally speaking, we cannot let privacy dependencies bear on the scope of *Excludability* in such a way to restrict third-parties in what information they may disclose about themselves or in general. This would presuppose that they can access information that are structurally unavailable to them.

The same point can be made if we focus on *false-negative errors*. Recall the genetic information case with which we began and assume that the brother decides not to share because he fears that others will infer his brothers' genetic information. As it turns out, the specific recipient *wouldn't have* inferred this (this could happen if they are unaware of the existence of the brother.) In this case, the brother who decided against sharing suffered the costs of a *false-negative error*, he let go of a benefit thinking his activity would violate his brothers' rights. Again, that the brother should bear such a cost doesn't seem acceptable because from their perspective they wouldn't be able to tell if the act would have been permissible or impermissible because it turns upon facts that are inaccessible to them. The brother cannot know if the specific recipient under consideration will be able to connect the dots and infer the genetic status of the other brother.

We think this argument shows why people cannot have a duty to avoid disclosing things about themselves to avoid that some get inferential access to other people's privacy-relevant information. Having such a duty would turn upon facts that are principally inaccessible to putative duty-bearers. Accordingly, we think we should reject that *Privacy Dependencies* can be relevant to the scoping of *Excludability* in such a way to require that people do not share information about themselves.

A worry might be that this argument proves too much. Why not think that it is *Excludability* that is the problem? Another way to put this worry is by questioning if people ever have access to the facts relevant for respecting *Excludability*. To address this concern, consider common examples of privacy violations: like reading someone's diary without their permission, secretly watching others, or sharing private information with third parties. In many of these cases, what those responsible for respecting privacy rights need to understand to avoid making errors is how information is shared through certain actions (such as reading a diary) and who is likely to receive this information (like the person reading the diary). And while duty-bearers will sometimes have false beliefs about these matters, it seems to us that people are in general quite well-attuned to the relevant facts. Humans are generally quite good at picking up on how information flows in their environment (compare Nissenbaum 2010).

Another concern may be that our conclusion is too strong. An anonymous reviewer raises the interesting challenge that our view implies that database owners wouldn't have a duty to minimize risks of de-anonymization when releasing anonymized datasets. But most tend to think, intuitively, that database owners have such duties and have them precisely for the sake of protecting the privacy rights of those people whose information is stored in the anonymized database. To respond to this challenge, recall that according to our argument, the scope of *Excludability* is a function of whether duty-bearers can reliably become aware of the negative consequences of their information-sharing activities. On this view, you are absolved from a duty

only if you have no reasonably available means to be in a good epistemic position with regards to the content of your duties. But it would seem to be the case that database owners in general have reasonable opportunities for becoming aware of, and subsequently mitigate, many risks that would otherwise ensue because of weakly anonymized databases. Partly because these risks are and should be well-known to professionals in the field, finding expression in things such as best-practice data protection standards and records of prior de-anonymization attacks. Thus, it seems to us that this challenge can be accommodated by paying attention to the mechanism that we suggest shapes the scope of *Excludability*.

III.B. The Waivability Dilemma

In the former section we saw that if we interpret *Excludability* broadly, it systematically puts duty-bearers in a bad epistemic position to comply with their resulting duty. This is an unattractive implication as duty-bearers thereby incur a serious risk of wronging others through their otherwise mundane information-sharing activities as they lack the evidence necessary to reliably avoid sharing information that can be used to infer something about third-parties. In this section we shall instead focus our attention on *Waivability* and whether it should be interpreted broadly in light of *Privacy Dependency*. We show that this exercise reveals a dilemma that threatens the idea that *Privacy Dependency* could motivate broadening the scope of *Excludability* much in the first place.

Suppose—contrary to what we argued in the former section—that *Excludability* should be scoped in light of *Privacy Dependency* in such a way that people can wrong others by sharing information about themselves that can be used in inferences targeting others. A natural question arises. Should we also broaden the scope of *Waivability*—the other central moral entitlement there is in the right to privacy—in light of *Privacy Dependency*? We suspect an affirmative answer to this question will be hard to avoid, but let's first see what it could mean to broaden the scope of *Waivability* in light *Privacy Dependency*.

Waivability says, roughly, that there is an entitlement to create permissions for others to access one's information. As we saw before, this can happen via sharing information but it can also happen in other ways. We can distinguish two views:

> **Waivability (narrow)**: when you create a permission to access a piece of personal information, *p*, about yourself, you create only this permission.
>
> **Waivability (broad**): When you create a permission to access a piece of personal information, *p*, about yourself, you create this permission, but also the permission to inferentially access any information that can be inferred from (or based partly on) *p*.

Consider an example to appreciate the distinction. Suppose that I send you a picture of myself at the beach taken during my time off. I plausibly create a permission (through waiver) for you to access this information by way of my act. But let's imagine that you can infer several other things from this picture. You can infer from the picture, in combination with background knowledge, that I have an early stage of skin cancer from the tone of my skin as depicted in the picture. You can also infer that my partner and I probably had a major conflict during the holiday since I am not wearing my wedding ring in the picture. In other words, things I might not want to share.

The narrow interpretation of *Waivability* would say that by sharing the picture I only create a permission to access the information contained in the picture. The broad interpretation would say that the created permission extends, as it were—to cover things I can infer, which I may access inferentially.[18] Notice that the broad interpretation need not say that the permission to infer things from the picture is a consent-based permission of a similar type to the permission created to access a photo. It could be that when I consent to somebody accessing a piece of information, I *forfeit* my right to not having this piece of information used in reasoning by others.[19] Or it could be that there simply exists no, or very few, prohibitions on using

[18] Rumbold and Wilson (2019) reject this; we'll discuss their view below.
[19] See Hanin (2022)

information (of a certain quality) you have available in your reasoning in the first place.[20] We need not settle on a view here, though.

Now, the dilemma. It follows from having to accept either the narrow or the broad interpretation of *Waivability*. Let's consider the narrow interpretation first. The problem with accepting this interpretation—at least on a background where we have accepted that *Excludability* should be interpreted broadly—is that it both seems *arbitrary* and *unfair*. It seems arbitrary because of the following explanatory burden: Why is it that *Excludability* must extend to protect inferential access, but *Waivability* must *not* extend to permit inferential access? Such an asymmetry calls for an explanation.

Unfairness next. To motivate this claim, we need to think about how the narrow and broad interpretation of *Waivability* affects the costs of error absorbed by the right-holder. Suppose that Milton and Miriam are close friends and that Milton tells Miriam where he spent his holiday last summer (and thereby waives the duty that Miriam would otherwise have to not access this information). Milton can't really tell what Miriam can infer about him based on this information—because this is determined by facts that are internal to Miriam—but it happens to be the case that she can infer a number of things that Milton wouldn't want to share. From Milton's perspective, then, the case in which he shares this information and Miriam will infer something he wouldn't want and the case in which she won't (or can't), looks identical as he cannot access Miriam's internal states that determine her inferential capabilities.

If we accept the broad interpretation of *Waivability,* there is a serious risk of Milton suffering the consequences of a *false-positive error*. This is because he cannot tell what Miriam can infer *in combination with* this interpretation of *Waivability* implying that he does not only waive his

[20]Some (perhaps) exceptions to this idea come to mind. For instance, perhaps certain principles of anti-discrimination implies that you cannot use information about sex, race or gender as a basis for making up your mind about certain matters. But it should be clear enough that this is a special case. See Lippert-Rasmussen (2013) for discussion of discrimination.

right to the information he shares, but also what others can infer from it based on the shared information. If he waives his right, and Miriam infers something he wouldn't want inferred (but couldn't reasonably foresee), there is a false-positive error and corresponding cost as there is a mismatch between what he intended to permit and what was actually permitted (The cost is of course that Miriam ends up doing something Milton wouldn't want).[21] However, if we accept the narrow interpretation of *Waivability* Milton will only ever waive his right to the information shared (not what can be inferred from it). Should Miriam infer something under the narrow interpretation of *Waivability*, she would bear the rectificatory costs for doing so (in virtue of becoming a wrongdoer) and should it be unclear to Miriam if she is allowed to infer something, she would bear the costs of corresponding errors.

It seems therefore that we should accept the narrow interpretation of *Waivability* in order not to let the uncertainty disadvantage Milton excessively. However, *that* judgment is hard to square with being indifferent to the costs imposed on putative duty-bearers required not to share their own information to not endanger the privacy of others as a result of their inability to predict how their information can be used in inferences targeting others. We can motivate this point as a case of interpersonal justification. Suppose a person who is bound by duties due to *Excludability* (in its broad interpretation) addresses a person who enjoys *Waivability* (narrowly interpreted). They can say something like:

> "*I* must take extreme care when I share information about myself in order to not share something that can be used to infer something about you. And *I'm* running a serious risk of wronging you because I am mostly clueless about what others can use my information to infer about you. So why shouldn't *you*—when you exercise your entitlement to waive and create permissions to access information—own the consequences of your conduct by

[21] We can imagine analogous cases where Milton will suffer a false-negative cost as result of not sharing something where he fears it can be used to infer something he wouldn't want to have inferred about himself and this is in fact not true.

rendering it permissible to access information (inferentially) about you from the information you've shared and thus absorb the same risk that I'm bearing?"

In other words, it seems unfair that duty-bearers should be held liable for what can be inferred from what they share (about others), whereas right-holders should *not* be held liable for what can be inferred (about them) from what they share.[22] Accepting this would amount to accepting that duty-bearers should unilaterally absorb the costs of error that result from it being unclear what a recipient can infer from a given piece of information. That strikes us as unacceptable, and this is the unwelcome implication of accepting the first horn of the dilemma that *Waivability* should be interpreted narrowly (in combination with saying that *Excludability* should be interpreted broadly). This strikes us as a serious challenge; it seems to us that the burdens and benefits are simply allocated unfairly and disproportionately benefit right-holders and disadvantage duty-bearers.

All this might push us towards the broad interpretation of *Waivability*, at least if we still are unwilling to say that people should be permitted to share information about themselves even when this can be used in inferences targeting others. What's the problem here? The problem here is entirely analogous to the one we motivated in the context of criticizing a broad interpretation of *Excludability* and what we encountered in the case of Milton and Miriam. If we demand of right-holders that they are held liable for what can be inferred from what they share, they will often end up being liable for things they simply had no reasonable way of anticipating. This is because they too will be in a bad epistemic position to determine what others can infer about them based on the information they share.[23] We saw that in the example with Miriam and Milton, and how it felt inappropriate to let Milton bear the costs of error from saying that when

[22] It might sound odd to be "held liable" for exercising a waiver. The idea would be that if a waivee is held liable for the information (about them) that can be inferentially accessed from what they disclose, then others would not be wronging the waivee by accessing the information inferentially.

[23] This point is reminiscent of Rumbold and Wilson's argument (2019) saying that you can't waive your right unintentionally, and therefore you cannot have said to waive your right over things that others could infer from what you shared.

he waives his right to a piece of information when sharing it, he also waives a right to what can be inferred from it.

Let's summarize our challenge. It turns out that scoping the right to privacy in light of privacy dependencies makes it extremely complicated for both right-bearers and duty-bearers to act informedly because it is mostly unclear what others can infer from a piece of information. We have argued that accepting a broad version of *Excludability* is objectionable on independent grounds, that it is objectionable to accept a narrow interpretation of *Waivability* in combination with a broad interpretation of *Excludability* (due to distributive unfairness), and that it is objectionable to accept a broad interpretation of *Waivability* (independently of what interpretation of *Excludability* you endorse). The only viable option, as far as we can tell, is accepting a narrow interpretation of both entitlements. This means, on the one hand, that people cannot be said to be wronging others when they disclose information about themselves that can be used to infer information about others. It also means, on the other hand, that right-holders cannot be said to waive their rights against inferences (if any such rights exist; we haven't argued for that conclusion here) when they share information about themselves.

	Waivability (narrow)	**Waivability (broad)**
Excludability (narrow)	Acceptable distribution of costs.	Unfair costs imposed on right-holders.
Excludability (broad)	Unfair costs imposed on duty-bearers.	Unfair costs imposed on both right-holders and duty-bearers

IV. Objections

We've been arguing that *Excludability and Waivability* are best interpreted narrowly, that is, in a way that disregards inferential privacy dependencies. The central idea we have been exploiting in

making this argument is that scoping these incidents in ways that take into account inferential privacy dependencies will result in a scoping that right-bearers and duty-bearers cannot navigate because they will typically not have structural access to information about what others might infer based on the information they share. This strikes us as a strong argument because of its generality, but it might also seem worrisome because the epistemic pattern meant doing work could admit exceptions. What should we say in such cases?

The most pressing challenge is of course saying what makes the genetic case with which we began special. But we think we are in a good position to explain that now. Genetic information is special, at least in part, because it is *common knowledge* that by sharing your genetic information you will—at least when information about how people are biologically related is easily publicly available—inevitably provide other people with all they need to infer parts of, or the entire genetic make-up of other people.[24] In cases of common knowledge, following Lewis (1969), it is not only true that two or more people know that P. For the knowledge to be common, each of them must also know (or rationally believe) that the other knows that P. Due to this latter property, common knowledge serves an important coordinating role in social life (for instance: car traffic would be much more difficult if we didn't have or couldn't assume that the traffic rules are common knowledge amongst trafficants). Accordingly, if it is common knowledge that a person's genetic information can be used to infer the genetic information of this person's biological relatives, then both information-sharers and information-recipients know this (or perhaps ought to know this) and they both know (or ought to know this). But if this is common knowledge, then we cannot say that this is something that people sharing (genetic) information about themselves couldn't be aware of.[25] Indeed, assuming that there is common knowledge about the inferential potential of genetic information seems to be in part what makes

[24] For discussion of the concept of common knowledge, see Vanderschraaf and Sillari (2023).

[25] This point can also be used to make better sense of the database case we discussed in footnote 18, although the knowledge needed in that need not be fully common. It seems to be enough that database owners can know what the inferential capabilities are of those who would attempt to de-anonymize the database; the recipients of the information needn't know that the database owners know of their inferential capabilities.

it plausible to think that Peter could be violating Carl's privacy in the case with which we began.[26] But we can also see that the common knowledge here is something that makes the genetic case very special. We're simply not attuned to the many other inferences people could be making (about us or others) based on the information we share, and it is unlikely that we will be except in a select few cases. Hence, the challenges we have mounted against accounts of the right to privacy that take *Privacy Dependency* into account can also be shown to point out why the genetic case is special.

Here is another worry about scope. One might fear that our argument speaks well to the privacy-dynamics in interpersonal encounters. But it may speak less well to the case of contributing one's information to the training of algorithms that will be deployed online to predict sensitive information about people (such as, for instance, those deployed by Google and Twitter). This is because, although we do not know what they will use our information for in particular (e.g., what they will be able to learn), we can justifiably believe that they are deploying towards such ends. To motivate this point further, notice that we often do not need to know the precise details of others' wrongful plans in order to have a duty not to contribute to them; it is enough that we know that somebody is up to something bad for us to have a duty not to contribute. Our response to this is that people might be doing wrong when they knowingly contribute their information to Big Tech and this information is used to train (and subsequently deploy) privacy-invading algorithms. We do not want to deny that it can be wrongful to contribute to wrongdoing. But we deny that this wrong has the distinctive profile of wrongs that are violations of the right to privacy. Remember that the paper started with an analogy suggesting that there is a relevant sense in which we are (wrongfully) sharing information about others when we share information about ourselves. This is an attempt to analogize the (normally regarded

[26] This point yields a recipe for predicting when we shouldn't expect people to have a duty to take into consideration privacy dependencies when sharing information about themselves: Namely, in cases where there is no common knowledge about what can be inferred from a given piece of information. In practice, then, an information-sharer must duly reflect upon the inferential capabilities of the information-recipients as common knowledge is something that is shared within a specific group of people.

permissible) case of sharing information about oneself with the paradigmatic case of (wrongfully) sharing information about others (as is the paradigmatic example of a violation of privacy). We have argued that this analogy should be rejected, but it doesn't follow from this that it is permissible to contribute one's information if it plays an important causal role in contributing to the wrongdoing of others. In fact, in the example with which we began we made no assumption about whether the person making the inference would be a wrongdoer.

A second objection to our argument could point out that the epistemic problem we have pointed out might not be insurmountable. To motivate this objection, imagine that our online environment was heavily transformed in the following way: Companies that live off making predictions about people (social media, search engines, etc.) suddenly began offering very clear notices on how they are using the information they collect about you (what things they are using it to infer about you as well as others) such that you could contract with them in ways that didn't put the privacy of others under threat. By stating these things up front, we could say that people could reliably become aware of how sharing their information could impact the privacy of others.

Our response to this is that we don't see that it should *in principle* be impossible to make such reforms (contracting on valid grounds can indeed shape the normative landscape in many ways). But we want to point out that they (in the best case) would only mitigate the epistemic problem in digital contexts, as it is hard to see how people could be required to disclose to others what they can or will likely infer from information they acquire. But aside from that, we are happy to grant that this might be the way forward for creating a more just online environment. But let us note that this scenario just seems extremely far away from where we are currently and that many are entirely skeptical of the idea that contracting - for instance via notice-and-consent - could do the necessary work (see for instance Nissenbaum and Barocas 2016). Hence, we suspect that such cases, while in principle imaginable, will be just as contrived and with little relevance for actual moral practice as the case with which our investigation began. So upon

closer inspection, *Privacy Dependencies* might still be mostly irrelevant to the scope of the right to privacy, at least in terms of how it affects permissions to share information about oneself.

V. Conclusion

Recall the example of Carl and Peter with which we began. We have argued that if there is a central moral insight to be gained from this example about the nature of privacy rights and the moral significance of privacy dependencies, then it is not an insight that generalizes beyond the rather unique case of genetic information. The reason, as we have suggested, is that duties that would result would take a form that duty-bearers couldn't be asked to comply with. What is left, we suggest, is that privacy dependencies are often irrelevant to what you are permitted to share about yourself. This, of course, doesn't rule out that privacy dependencies are morally significant in other ways. We'll mention two such possibilities in closing.[27] First, privacy dependencies may be morally significant in the sense that people who routinely seek to exploit them - as opposed to people who contribute the information needed to exploit them - may sometimes have duties to not do so. This thought chimes well with an argument made by Rumbold and Wilson (2019) according to whom companies engaging in predictive analytics may have special duties to respect the privacy of those that are the target of such predictions. Second, privacy dependencies may give rise to collective duties. According to such a view, privacy dependencies can be thought of as a negative externality, and sometimes the best response to externalities is to rely on the state to enact and enforce regulation that puts a ban on unreasonably risky activities. And perhaps we owe it to each other - as a collective - to create regulation (e.g., strengthen privacy law) that limits the negative effects of privacy dependencies. Interestingly, if we are right in saying that privacy dependencies typically don't affect people's permissions to share information about themselves, then this may be taken as some indirect evidence in favor of these suggestions.

[27] Thanks to an anonymous reviewer for prompting us to discuss these points.

References

Barocas, Solon, and Karen Levy. "Privacy Dependencies." *Washington Law Review* 95.

Barocas, Solon, and Helen Nissenbaum. "Big Data's End Run around Anonymity and Consent." In *Privacy, Big Data, and the Public Good*, edited by Julia Lane, Victoria Stodden, Stefan Bender, and Helen Nissenbaum, 1st ed., 44–75. Cambridge University Press, 2014. https://doi.org/10.1017/CBO9781107590205.004.

Bolinger, Renée Jorgensen. "Moral Risk and Communicating Consent." *Philosophy & Public Affairs* 47, no. 2 (2019): 179–207. https://doi.org/10.1111/papa.12144.

———. "The Moral Grounds of Reasonably Mistaken Self-Defense." *Philosophy and Phenomenological Research* 103, no. 1 (2021): 140–56. https://doi.org/10.1111/phpr.12705.

Floridi, Luciano ; Taylor, Linnet & van der Sloot, Bart (eds.) (2017). Group Privacy: New Challenges of Data Technologies. Cham: Imprint: Springer.

Fried, Charles. "Privacy." *The Yale Law Journal* 77, no. 3 (1968): 475–93. https://doi.org/10.2307/794941.

Gebru, Timnit, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, and Li Fei-Fei. "Using Deep Learning and Google Street View to Estimate the Demographic Makeup of Neighborhoods across the United States." *Proceedings of the National Academy of Sciences* 114, no. 50 (December 12, 2017): 13108–13. https://doi.org/10.1073/pnas.1700035114.

Hanin, Mark. "Privacy Rights Forfeiture." *Journal of Ethics and Social Philosophy* 22, no. 2 (July 26, 2022). https://doi.org/10.26556/jesp.v22i2.1633.

Inness, Julie C. (1992). *Privacy, Intimacy, and Isolation*. New York, US: OUP.

Kevin Macnish, Mass Surveillance: A Private Affair? - PhilPapers." Accessed March 13, 2023. https://philpapers.org/rec/MACMSA-12.

Lewis, David, 1969, *Convention: A Philosophical Study*, Cambridge, MA: Harvard University Press.

Lippert-Rasmussen, Kasper (2013). *Born Free and Equal? A philosophical inquiry into the nature of discrimination.* Oxford: Oxford University Press.

Lundgren, B. A Dilemma for Privacy as Control. *J Ethics* 24, 165–175 (2020)

Mainz, Jakob. "An Indirect Argument for the Access Theory of Privacy." *Res Publica* 27, no. 3 (2021): 309–28. https://doi.org/10.1007/s11158-021-09521-4.

———. "Inferences and the Right to Privacy." *Journal of Value Inquiry*, n.d., 1–19. https://doi.org/10.1007/s10790-022-09911-8.

Mainz, Jakob Thrane, and Rasmus Uhrenfeldt. "Too Much Info: Data Surveillance and Reasons to Favor the Control Account of the Right to Privacy." *Res Publica* 27, no. 2 (2021): 287–302. https://doi.org/10.1007/s11158-020-09473-1.

Marmor, Andrei. "What Is the Right to Privacy?" *Philosophy & Public Affairs* 43, no. 1 (2015): 3–26. https://doi.org/10.1111/papa.12040.

Menges, L. A Defense of Privacy as Control. *J Ethics* 25, 385–402 (2021)

Moore, Adam D. "Privacy, Interests, and Inalienable Rights." *Moral Philosophy and Politics* 5, no. 2 (2018): 327–55. https://doi.org/10.1515/mopp-2018-0016.

Munch, Lauritz, and Jakob Mainz. "To Believe, or Not to Believe – That Is Not the (Only) Question: The Hybrid View of Privacy." *The Journal of Ethics*, January 26, 2023. https://doi.org/10.1007/s10892-023-09419-8.

Nissenbaum, H., 2010, *Privacy in Context: Technology, Policy, and the Integrity of Social Life*, Stanford: Stanford University Press

Owens, David. *Shaping the Normative Landscape*. Oxford, New York: Oxford University Press, 2014.

Rachels, James. "Why Privacy Is Important." *Philosophy & Public Affairs*, vol. 4, no. 4, 1975, pp. 323–33.

Raz, Joseph. *The Morality of Freedom*.. Oxford University Press, 1986.

Rumbold, Benedict, and James Wilson. "Privacy Rights and Public Information." *Journal of Political Philosophy* 27, no. 1 (2019): 3–25.

Rössler, Beate (2004). "The Value of Privacy". Polity Press.

Thomson, Judith Jarvis. "The Right to Privacy." *Philosophy & Public Affairs* 4, no. 4 (1975): 295–314.

Véliz, Carissa. *Privacy Is Power*. London, UK: Penguin (Bantam Press), 2020.

Vanderschraaf, Peter and Giacomo Sillari, "Common Knowledge", *The Stanford Encyclopedia of Philosophy* (Winter 2023 Edition), Edward N. Zalta & Uri Nodelman (eds.), forthcoming URL = <https://plato.stanford.edu/archives/win2023/entries/common-knowledge/>.

Wang, Yilun, and Michal Kosinski. "Deep Neural Networks Are More Accurate than Humans at Detecting Sexual Orientation from Facial Images." PsyArXiv, September 7, 2017. https://doi.org/10.31234/osf.io/hv28a.

Westin, Alan. "Privacy And Freedom." *Washington and Lee Law Review* 25, no. 1 (March 1, 1968): 166.