The Epistemology of Privacy

MSc Thesis *(Afstudeerscriptie)*

written by

Laura Hernández

under the supervision of **Dr. Tom Schoonen**, and submitted to the Examinations Board in partial fulfillment of the requirements for the degree of

MSc in Logic

at the *Universiteit van Amsterdam.*

Date of the public defense:	**Members of the Thesis Committee:**
February 20, 2025	Dr. Marjolein Lanzing
	Dr. Aybüke Özgün
	Dr. Tom Schoonen
	Dr. Petter Törnberg (*chair*)

Institute for Logic, Language and Computation

Abstract

Privacy is ever so important in the digital age. To demand protection against privacy invasions, a necessary preliminary step is to elucidate what privacy is. This thesis aims to investigate the nature of privacy via two main research directions. First, throughout this thesis, I motivate and defend that our intuitions about privacy respond to two separate notions: descriptive privacy and the normative right to privacy. For this purpose, I present, defend, and expand on the Hybrid Account of Privacy by Véliz (2024). Second, I show that insights from epistemology can aid in understanding and formulating more clearly the descriptive and normative definitions of privacy.

On the one hand, for the epistemology of descriptive privacy, defined as a ternary relational state, I discuss relevant positions in the literature to study the epistemological nature of (descriptive) privacy losses. I argue that subject A loses privacy over their personal information p with respect to individual B if three conditions are satisfied: (1) p is true, (2) B believes that p, and (3) B's true belief is linked by some form of epistemic merit, i.e., it is justified, formed in a reliable process, or has a causal link.

On the other hand, for the epistemology of the normative right to privacy, I present, defend, and expand on the right to robust privacy by Véliz (ibid.), where a subject has a right against privacy invasions in the actual world and counterfactual worlds. Then, I embed this right to robust privacy within an epistemic pathway that analyses the flow of personal information p in three steps: (1) the inquiry, (2) the access, and (3) the use of p. I propose to reformulate a condition protecting the right to robust privacy in step (1) by using the epistemological notion of *knowability*. Finally, I discuss further research directions to expand on the investigation from this thesis.

Acknowledgements

I would first like to thank my supervisor Tom Schoonen. Thank you for believing in this research project and for your encouragement, support, and help along the way. Writing this thesis in five months was challenging, and I truly believe it was only possible thanks to your guidance, determination, and our exchange of thoughts and ideas—especially when our meetings extended for more than two hours. Thank you for grounding me and helping me focus my thesis—this document did not turn into (only) a compilation of very disperse ideas thanks to you!

I would also like to thank my defense committee, Aybüke Özgün and Marjolein Lanzing, for taking the time to read my thesis and for their thought-provoking questions during the defense. Thank you for your inspiring modules in *epistemology* and the *philosophy of technology*, which largely motivated this thesis. Thank you to Petter Törnberg for chairing the defense and, in that note, thank you to everyone who attended my thesis defense and shared a very special moment with me. I would also like to thank my academic mentor, Dick de Jongh, for his guidance, mentoring, and advice throughout the master's.

I was lucky enough that the topic of my thesis was built on work throughout my entire master's. Being able to study mathematical logic, formal epistemology, natural language processing, and different kinds of philosophy in the same master's made this combination fruitful. Most notably, this thesis would not have been possible without the support and encouragement I received during the second year of my master's, when I became incredibly determined to research the interplay between privacy and technology in every module I followed. Thank you to all of my professors and fellow students for their inspirational ideas during lectures, philosophical discussions, and essay writing. I have not only learned about philosophy, but I have also learned a lot about life and what makes it valuable. Your relentless insights will continue to guide my fights for privacy, democracy, and fairness.

Thank you to my friends in the master's, the *illogicians*, for sharing laughs, contradictions, and nonsense. We took a postgraduate degree to finally learn some common sense and I am unsure we succeeded. I have you in my heart.

In that note, thank you to every single person who made Ámsterdam (and Mánchester) feel like home. For privacy reasons (and for fearing I will forget someone), I will leave the suspense by not listing people. To all of you, I look forward to meeting each other in the future.

Lastly, thank you to my family; without your support, way of being, and never-ending love, I would not be the person that I am, and I would have never left home to run around the world and learn about such exciting things. Thank you for encouraging—especially—my most ludicrous dreams. Os quiero muchísimo.

Contents

Introduction

We live in unprecedented times. The power that big technological companies exert over our data and our actions increases each day. In the digital age, to live a standard life, we are tied to using digital platforms mediated by technology monopolists. We use them to stay in contact with our friends, share our creative work, communicate with family members, create documents and presentations for work, and to look for housing or sell our furniture. Meanwhile, big tech creates and collects more information about us each day, which they store in water-hungry and energy-devouring data centres that take over our local villages (Schaake 2024, pp. 59–66). It is not only the data created via our interaction on their platforms that they collect, anything published on the internet, from journalists' pieces to musicians' songs, is scraped by web crawlers to build generative artificial intelligence (AI) systems that they force us to use. Generative AI systems that they market to us as *automated* and *independent*, rendering invisible the painstaking traumatic work from data labellers, data annotators, and content moderators, whose work is crucial to building them (Williams 2024). Further, worsening the illness then forcing an ineffective remedy drug upon us, they promise to save us from the disasters of climate change by building AI systems that use more water and energy than many households in a year (Bender et al. 2021, pp. 612–613; Minde 2023). Meanwhile, people continue being unfairly denied loans, accused of fraud and crimes they did not commit (GOV.UK 2024; Heikkilä 2022; Murray et al. 2023), and being denied entrance to countries (Chandler 2025) thanks to *predictive* technologies that, in reality, do not predict but rather *dictate* our future (Susser et al. 2019). Tech monopolists have more money than most democratic states, and they are now beginning to overtly occupy positions of power in governments (Schleifer and Ngo 2025). The state of human rights, democracy, and education in our digital age are uncertain, with these issues being connected in an intricate web of interests, injustices, and inequalities. While the poor become poorer and the rich become richer, how do we navigate the challenges arising from our current digital world? In this thesis, I do not solve all of the issues raised. Instead, I aim to develop our understanding of a value that has a connection to these issues and that is becoming increasingly critical in this digital age: *privacy*.

Protecting our privacy can, in turn, protect us from abuses of power (Véliz 2020; Véliz 2024). For years, an imbalanced power dynamic has emerged between what big technological companies know about us and what we know about them. As it becomes more difficult to live a normal day without interacting with a device that is connected to the internet, the instants of our lives that are not collected by technology companies decrease. Many scholars, not-for-profit organisations, journalists, and human rights activists warn us that our dependency to digital devices poses a threat to our privacy, but what is *privacy*?

To begin to protect our privacy, and to demand for our right to privacy to be respected, one might begin by clarifying what privacy is. This thesis aims to do exactly that. This thesis aims to provide a foundational investigation on the nature of privacy by using insights from the field of epistemology. For this purpose, I motivate and defend that our intuitions regarding privacy respond to the necessity of separating privacy into two notions: *descriptive* privacy and the *normative* right to privacy. The former concept captures our morally neutral intuitions regarding privacy, while

the latter describes the inherently moral right to privacy. I discuss these two notions in detail throughout the thesis and I investigate their epistemology in Chapters 2 and 3, respectively.

The main **content** of this thesis is divided into three chapters. In Chapter 1, Section 1.1 motivates with examples the idea that our intuitions regarding privacy should be divided into a descriptive and a normative notion. To define these notions, I present my formulation of the Hybrid Account of Privacy by Véliz (2024). I defend Véliz's account by surveying and discussing the main access and control accounts of the philosophy of privacy literature in Section 1.2. I show that access accounts are best suited to respond to our descriptive conception of privacy, while control accounts better correspond to our normative intuitions. I further argue for this differentiation in Section 1.3.

In Chapter 2, I provide an epistemological analysis of (informational) descriptive privacy. For this purpose, first, I define descriptive privacy as a ternary relational state following an analysis by Blaauw (2013). I expand on Blaauw's analysis by investigating the three relata of the privacy relation in Section 2.1, i.e., the subject A, who is the being whose privacy is under evaluation, the set of personal information P about A, and the set of individuals I, who are capable of accessing P. After analysing the relational state of privacy, I turn to discussing the phenomena of (informational) privacy losses in Section 2.2, which were argued to be described in terms of *epistemic access* in the previous chapter. I argue that B epistemically accesses a piece of personal information p about A (leading to A losing privacy to B regarding p) if the following three conditions are satisfied: p is true (Section 2.3), B believes that p (Section 2.4), and B's true belief is linked by some form of *epistemic merit*, e.g., justification, reliability, or causality (Section 2.5).

The epistemology of the normative right to privacy is discussed in Chapter 3. As a first step, I present and defend the right to robust privacy by Véliz (2024) in Section 3.1. Then, I defend Véliz's right to privacy against recent objections from Munch (forthcoming). Subsequently, I argue that Véliz's right to privacy fails to account for misuses of personal information, and I propose a way to embed it into her conditions. Then, in Section 3.2, I motivate the normative analysis of privacy in terms of the pathway surrounding the act of accessing someone's personal information. In this sense, I argue for the pre-access step of the epistemic pathway to be described by the action of inquiry and I propose to reformulate the conditions that protect this step in terms of the epistemological notion of *knowability*. I briefly discuss some considerations given in the post-access step, especially its necessity to protect us from abuses of power that are detriment to our autonomy.

The main **contributions** of this thesis to the research fields of privacy and epistemology, and their intersection, are the following. First, I provide an explicit argument for considering privacy as having two distinct natures (descriptive and normative), responding to a long-standing discussion on the nature of privacy. Then, I discuss notions in epistemology that help in the analysis of the two notions of privacy - a strategy that has hitherto been reserved for descriptive privacy. For the epistemology of descriptive privacy, I provide a novel objection against an argument from Véliz determining that a *lucky true belief* can yield a privacy loss as long as its consequences are undistinguishable to those caused by a *justified true belief*. I argue that Véliz's argument excessively relies on pragmatic consequences that confuse the effort of descriptive privacy with that of normative privacy. For the epistemology of the normative right to privacy, I provide a novel defence of Véliz's right to robust privacy against recent objections raised by Munch (ibid.). I further notice that Véliz's right to robust privacy fails to account for misuses of personal information as right to privacy violations, and I propose ways for Véliz to amend this within her theory. Lastly, I tentatively propose an epistemic pathway as a manner to epistemologically analyse the right to (informational) privacy more holistically. In the Conclusion, I discuss a more detailed list of smaller contributions from this thesis and further work to build on the investigations provided here. For now, let us begin by discussing the notions of descriptive and normative privacy in Chapter 1.

Chapter 1

Descriptive and Normative Privacy

To investigate the nature of privacy, we can begin by examining processes which entail a loss of privacy, and so build up a definition of what it means to have privacy at all. In justifying this starting point, try to determine how protected your privacy is currently. You may remember the last privacy policy you accepted without reading it to subscribe to some service on the internet.[1] You may look around to check for any cameras in the room or for people at eyesight distance from your computer screen. If you are in a public space, you may also notice the clothes you are wearing to cover some parts of your body. We judge how protected our privacy is at a particular time by entertaining cases in which it could have been lost by, e.g., a digital company misusing our personal data or strangers watching us.

In this chapter, I will investigate the nature of privacy by examining these privacy diminishing cases and finding what it is that unifies them. In order to provide a working definition for privacy, I will determine what (lack of) actions ensure its protection. My main aim will be to clearly show that privacy is a two-fold notion that has a descriptive and normative side. For this purpose, in Section 1.1, I will explore some preliminary examples to spark intuitions on the descriptive and normative notions of privacy. Then, I will present and motivate the hybrid account of privacy by Carissa Véliz (2024), where the descriptive aspect of privacy is defined as lack of access, while the normative aspect is conceptualised as a right to robust privacy.[2] To argue for this hybrid account, in Section 1.2 I will present and discuss the main theories of privacy in the literature: the access accounts and the control accounts. In Section 1.3, I will further clarify and present the descriptive and normative notions of privacy that I am endorsing in this thesis and I will argue why it is important to distinguish between them. Section 1.4 concludes this chapter.

1.1 The nature of privacy

1.1.1 Examples

Let us begin by presenting some examples to help us gain intuition on what privacy is. For the following, let A, B, and E be distinct individuals, and let p be a personal proposition about A.

[1] A North American study found that it would take an average of 244 hours (around 30,5 full working days) per year for the average North American to read every privacy policy they accept on the internet. These were estimated to be 1462 privacy policies per year, as calculated in 2008 (McDonald and Cranor 2008). Given the increment in internet use from 2008 until now, it seems safe to assume that the average time in 2025 has increased considerably.

[2] In Chapter 2, I will investigate the epistemological nature of the descriptive notion of privacy, conceptualised as a relational state. In Chapter 3, I will present and discuss Véliz's right to robust privacy more closely, and supplement it with an epistemic pathway for the flow of personal information.

Example 1.1.1 (**Café**). *A* is chatting with *B* in a café. *A* discloses personal information *p* to *B* willingly. *B* hears *p* adequately. No one eavesdrops on their conversation.

Does *A* lose privacy to *B* over *p*? Our intuitions might diverge, representing the main ideas from the literature. Some people argue that, in this example, *A* does not lose privacy because *A* deliberately chooses to share *p* with *B*. In doing so, *A* exercises *control* over their personal information. People defending this view highlight that we regularly find ourselves in situations where we disclose personal information, e.g., when we go to the doctor, we fill in a form for our town hall, or we chat with our partner. Due to the ordinary nature of these cases, they argue that it would be strange to conceive these as scenarios in which we lose privacy. This position requires that for an agent to have privacy, they must exercise some form of control over their personal information. In the literature, this view is called the *control account* of privacy.

In contrast, the *access account* conceives privacy as the lack of access to personal information. Notice that the access account does not require the person whose privacy is under assessment to exercise any agency to protect their privacy. Within an access theory, *A* loses privacy to *B* over *p* the moment that *B* hears *p* adequately simply because *p* is personal information.[3] However, losing privacy is not a bad thing *per se* as, according to this view, it is very common in our society to relinquish our privacy for purposes such as building or strengthening a personal relationship, for receiving necessary help from a doctor or a lawyer, or for being an active citizen in society.

So, according to the control account, *A* does not lose privacy over *p* with respect to *B* in Example 1.1.1, while *A* does lose privacy according to the access account. Let us consider a modified version of this example where our moral compass becomes more salient to continue testing our intuitions regarding privacy.

Example 1.1.2 (**Café and *E***). *A* is chatting with *B* in a café. *A* discloses personal information *p* to *B* willingly. *E* eavesdrops on their conversation. *B* and *E* hear *p* adequately.

According to the access account of privacy, *A* loses privacy over *p* to both *B* and *E* in this situation because they both hear *p* adequately. From this, one may notice a potential weakness in the access account, as both *B* and *E* are considered to equally diminish *A*'s privacy (since they both access *p*).[4] This way, the access account fails to capture the different level of permissibility between the actions that *B* and *E* perform to access *p*, i.e., this account fails to capture the normative intuitions of privacy. In contrast, within the control conception of privacy, *A* only loses privacy to *E* because *A* does not choose for *E* to listen *p*, while *A* does choose to disclose *p* to *B*. In fact, *A* is not even aware that *E* eavesdrops. The control account concludes that, by eavesdropping, *E* diminishes *A*'s privacy. To explore further whether a moral difference in *E*'s actions translates into a different status for privacy, consider the following two situations:

(1) *E* happens to be seating nearby *A* and *B*. *E* eavesdrops on *A* and *B*'s conversation with little effort,

(2) *E* eavesdrops on the conversation with the use of an advanced hearing device, which required planning and effort.

The main difference between these two situations pertains *E*'s intentions and the amount of (unjustified) effort that *E* utilises in eavesdropping. In (1), *E* may have acted wrongly by not ignoring

[3] One may wonder whether hearing *p* is enough for a privacy loss or whether *B* needs to acquire a strong epistemic state, such as knowledge. For instance, imagine that *B* is on their phone and hears *A* without paying attention. One may be sceptical that *A* loses any privacy if *B* is not attentive enough. For now, I assume that hearing personal information is sufficient for our present purposes. I will investigate the epistemic requirements for *accessing* personal information in the next chapter.

[4] It will become clear in the following chapter that, if one takes privacy to be a gradual notion, then one can distinguish between *B* and *E* within the access account by positing that *E* diminishes *A*'s privacy more than *B* since, for instance, *E* hears *p* without *A*'s awareness or consent.

A and B's conversation to the best of E's abilities. In fact, one may even consider it a moral duty for E, and any other person, to ignore a personal conversation whenever possible. However, E still gained this information circumstantially so, intuitively, the moral wrong from E's action is relatively small. In contrast, situation (2) is more wrongful precisely because of the intentions and the extra effort that E exercises in targeting A and B's conversation by using an advanced hearing device. Intuitively, this might mean that A loses more privacy over p in (2) than in (1) because A loses more *control* in (2) than in (1). In particular, one may highlight that E's exercising such effort to acquire p makes us especially uneasy as it seems more likely that E will use p for a further purpose, such as harming A. It is clear that, in this situation, E wrongs A by eavesdropping, regardless of what they do with p. However, it is also intuitive that E would wrong A further and, diminish their privacy more, were E to disclose p to other people. Therefore, our normative intuitions regarding privacy are not exclusive to the fact that a privacy diminishment occurs. Instead, moral considerations are also dependent on the entire situation surrounding the privacy diminishment, such as the manner that p is acquired and the use of p afterwards. I will discuss these moral considerations in Chapter 3.

1.1.2 Definitions

The previous examples aimed at challenging preconceived ideas on privacy by contrasting an ordinary situation where no wrong or harm occurs, even though privacy is arguably diminished (Example 1.1.1), with a case where an external agent eavesdrops on a conversation, thus wronging and (arguably) harming the person being eavesdropped by breaking social norms, and where privacy has been more clearly diminished (Example 1.1.2). At this point, before I discuss the access and control accounts of privacy further, it is important to differentiate between a *normative* and a *descriptive* notion of privacy. The main argument of this thesis is that privacy can be defined, on the one hand, in terms of the moral intuitions sparked in Example 1.1.2, which relate to the right to privacy and its value in, e.g., protecting us from being eavesdropped. This is the normative side of privacy. On the other hand, privacy can also de defined in morally neutral terms. This is the descriptive side of privacy and it responds to intuitions deeming that one's privacy can be lost even in a situation that sparks barely any moral reactions, such as in Example 1.1.1.

In this section, I will introduce, motivate, and define the main concepts that I will use throughout this thesis: personal matters, and the descriptive and normative notions of privacy.

Personal matters

Before we try to elucidate what privacy is, let us begin by inquiring on the *things* that we protect when we talk about preserving our privacy. I will refer to these as *personal matters*. In the following, I will motivate for these to be personal information and personal space. For this purpose, consider the following definition where A is an individual.

Definition 1.1.1 (Personal matter). p is a *personal matter* of A if

(i) p is a piece of personal information about A, or

(ii) p is personal space surrounding A.

Personal information and personal space are difficult terms to define. In particular, not every piece of information about someone that one learns causes a privacy diminishment, e.g., the fact that I have five fingers on each hand. Similarly, not the same amount of physical space that surrounds someone causes a privacy diminishment each time, e.g., compare the difference in the amount of space that we consider appropriate for others to invade when we are in a busy metro, as opposed to when we are in a public toilet. Context and culture partly influence what counts as

personal information and space. Emerging technologies provide an extra challenge for drawing the boundaries of these definitions, as I will exemplify below with the challenges that big data poses for personal information and virtual reality for personal space. Thus, it is not possible to delineate these notions precisely in a simple way and I will not attempt to do so in this thesis. Instead, I will provide some basic definitions for what I refer to when I talk about *personal information* and about *personal space*, and I will raise some issues regarding these definitions.

A first step towards a definition of **personal information** can be found in the European Union's General Data Protection Regulation, where *personal information* is defined as "any information relating to an identified or identifiable natural person" (European Parliament and Council of the European Union 2016, May 4, Art.4). From this definition, someone's national identity card, social security number, and their biometric information are all instances of that person's personal information because they relate to them. However, by relying on information related to someone, this definition encounters challenges in the age of big data because an increasing amount of information is accumulated, aggregated, and inferred from an individual, and seemingly innocuous data can now be related to a person. For instance, social media activity can be used to infer whether a user is depressed (De Choudhury et al. 2021; Salas-Zárate et al. 2022), and supermarket shopping habits to determine whether they are pregnant (Duhigg 2012). This is problematic for the question of what personal information is as the age of big data makes even the most trivial of information related to us relevant for privacy, thus turning it into personal information.

The notion of **personal space** is also problematic to define. To see this, let *personal space* refer to a varying region of perceptual and physical space that surrounds a person. This is an inescapably vague definition. Personal space has to account not only for physical space, but also for perception because, for instance, if someone stares at you for an unusual amount of time, even when located in an opposite side of a room, it can feel like they are invading your personal space. Personal space is also a varying notion because it will depend on our culture, preferences, and notion of comfort whether it will seem permissible or impermissible for someone to stand very close to us. It will also depend on our situation, as it is not the same to be very close to strangers while watching a concert than while having dinner. In itself, personal space presents many interesting challenges to our notion of privacy. The increased use of augmented and virtual reality also challenges its definition as we may suddenly feel like our personal space is exposed if what we see and listen is mediated by a corporate device. I will not enter into the intricacies of personal space in this thesis, as I will mainly focus on informational privacy (i.e., privacy with respect to personal information), although many interesting issues can be raised, and so personal space presents itself as a great avenue for further research.

Some philosophers in the legal and moral literature have considered for privacy to protect something distinct to our personal information and our personal space. *Decisional privacy* is the most prevalent of such efforts. Decisional privacy refers to the idea that for a subject's privacy to be protected, not only should external people lack epistemic access to certain information about the subject and lack perceptual and physical access to certain space, people should also refrain from intervening in the subject's personal decisions (Rössler 2005). Decisional privacy is specially relevant in the digital age as automated systems are built using our personal data and are deployed in apps and services that target advertisements to us and recommend us certain content, affecting our actions and decisions (Lanzing 2019). Now, when it comes to the relationship between *decisional privacy* and the analysis of privacy relates to a difficulty in differentiating descriptive privacy from the notion of autonomy (Véliz 2024, p. 56), since autonomy is usually related to some form of non-intervention.

In this thesis, I will consider decisional privacy as an important notion that protects a subset of what autonomy protects. In particular, decisional privacy protects the decisions that are under the threat of intervention by an external entity as a result of them accessing our personal information or

space. However, due to their inherent connection to morality, I will not embed decisions within the definition of *personal matters*, which refer to the privacy objects whose access must be protected to ensure descriptive privacy. Instead, I will conceptualise decisional privacy in the normative right to privacy. Specifically, I take decisional privacy to impose meaningful restrictions on the use of personal information within an epistemic pathway for the protection of the right to privacy. These ideas will become clearer once we have defined descriptive and normative privacy.

Descriptive privacy

When it comes to *descriptive* privacy, Ruth Gavison (1980) is one of the main figures stressing the importance of elucidating a neutral and descriptive notion of privacy. She notes that specifying a descriptive notion of what a diminishment of privacy amounts to, before we try to assess whether it is good or bad, can be instructive for inspecting the basic nature of privacy (ibid., pp. 423–425). In this thesis, I will argue for descriptive privacy to be defined within an access account. That is, one has descriptive privacy over a personal matter with respect to external people if those people lack access to one's personal matters. I will present this definition in this section and defend it in the next section against the so-called control account of privacy. For the following, let A and B be distinct individuals, and let p be a personal matter about A.

Definition 1.1.2 (**Descriptive Privacy**)**.** A has *descriptive privacy* with respect to B over p if and only if B lacks access to p.

Once we have a definition of descriptive privacy, we can define privacy losses (and diminishments) in the following way.[5]

Definition 1.1.3 (**Privacy loss/diminishment**)**.** A *loses privacy* to B over p (i.e., B *diminishes* A's *privacy* over p) if and only if B accesses p. Specifically, this means that:

1. If p is A's personal information, then B epistemically accesses p.

2 If p is A's personal space, then B perceptually (and phenomenologically) accesses p.

The distinction between epistemic and perceptual access relates to the variability of personal matters as personal information or personal space. In this thesis, I will only inspect the epistemic side of privacy and so will focus on privacy with respect to personal information. Now, within definition 1.1.3, a privacy loss or diminishment is not necessarily something bad. To see this, recall Example 1.1.1, where A and B are in a café and A discloses personal information p to B willingly. In this case, A has neither been wronged nor harmed, but A still loses privacy because B accesses p. This explains why privacy losses and diminishments are tied to *descriptive* privacy, which is a normatively neutral conception of privacy. The usefulness of defining descriptive privacy, in contrast to normative privacy, will become clearer in the following sections, when the access and control accounts are discussed. In short, descriptive privacy allows to signal a subject A's position of vulnerability with respect to another individual B, regardless of whether B exploits this vulnerability or not (and whether this leads to moral considerations).

Normative privacy

It is important to note that privacy is a heavily normative term and that the motivation for its philosophical debate stems mostly from instrumental reasons (Munch and Mainz 2023, p. 256). For instance, it is widely accepted in the privacy literature that an important function of privacy is for it to enable other goods and rights we consider valuable, such as our autonomy, well-being, creativity, democracy, and so on. In recent years, a surge of philosophical work has emerged

[5] I use privacy losses and privacy diminishments interchangeably to refer to the same phenomenon.

that defends the value of privacy with respect to the increased deployment and normalisation of surveillance technologies. The seminal work on philosophy of privacy, considered the first main philosophical work on the topic, was written by Warren and Brandeis (1890) as a defence of privacy against the threat of the indiscriminate use of instant photographs by reporters. In our current digital age, increasingly, *smart* vacuum cleaners, dishwashers, and other equipment with access to the internet, and mostly equipped with a camera and audio, are present in our homes, or in our friend's homes. Health tracking devices and apps monitor our every move. Even our smartphones and *smart* cars often have geolocation services turned on by default, allowing for the collection and potential dissemination of our precise location at all times. Even when our smartphones and the developers of the apps themselves are unaware of it, our location data is often covertly collected and disseminated by advertising infrastructure (Cox 2025a). This way, the offline and the online worlds merge as surveillance technologies become the norm.

In the digital age, a normative conception of privacy must be capable of critiquing surveillance practices because of the manner in which they collect and disseminate our information, and because of what this means to our values and human rights. For this reason, the normative conception of privacy that I will endorse does not only account for whether our privacy is diminished, but also for the wrong that its diminishment and even the threat of its diminishment can cause. It follows that the right to privacy is a *path-based* notion as we not only care about whether someone invades our personal information, but we also care about whether they attempt to access it, and the manner in which they access it. In particular, these considerations change how morally wrong the action is. In contrast, descriptive privacy is a *path-independent notion* because, to assess whether we have descriptive privacy or not, the only relevant thing to consider is whether a piece of personal information is accessed or not; it does not matter *how* it is accessed. My conceptualisation of the normative right to privacy is based off Véliz's right to robust privacy (Véliz 2024, p. 145). The right to robust privacy ensures that descriptive privacy is protected in the actual world, and also in **counterfactual** scenarios.

I will now present a definition of the normative right to privacy based on Véliz's theory. Note that Véliz does not provide an explicit positive definition of the right to robust privacy; therefore, the formulation I present here negates the actions that lead to a violation of the right to privacy as defined in (ibid., p. 145).[6] I will further differ in my presentation of the definition of the normative right to privacy by explicitly stating possible worlds as $w_1, w_2, ...$ to make a clear connection with possible world semantics. Note that Véliz does not use this formulation within her definition.

Let A and B be distinct individuals, and let p be the personal matters of A, let $w_1, ..., w_n$ be possible worlds, and let *invading A's privacy* correspond to *diminishing A's descriptive privacy without A's consent.*

Definition 1.1.4 (**Normative Right to Robust Privacy**). *B does **not** violate A's normative right to robust privacy* at w_1 if

(i) At w_1, B does not intend to invade A's privacy at w_1 or at relevant possible worlds $w_2, w_3, ...$ (i.e., this is a subset of the whole set of possible worlds).

(ii) At w_1, B does **not** attempt to invade A's privacy, and

(iii) at w_1, B does **not** invade A's privacy.

Notice that the normative right to privacy is **robust** since it requires for the external person B to not attempt to invade A's privacy, to not have the intention to invade it in actual or counterfactual

[6]Note that Véliz also discuss two further notions regarding the right to privacy: infringements and failures to respect this right (Véliz 2024, pp. 145–146). I will present and discuss these notions in Chapter 3. For now, it is important to note that the act of violating the right to privacy is the most morally condemnable action related to the right to robust privacy, and this is why I portray the right to robust privacy as a negation to the violation of this right. In that sense, realising the lack of actions necessary to avoid the violation of someone's right to privacy will suffice to show the *robust* nature of Véliz's formulation of the right to privacy.

worlds, and to not actually invade it. It is important to note that Véliz considers her right to robust privacy as distinct to a control account of privacy. Recall that control accounts consider someone's privacy to be respected as long as the person has (some form of) control over their personal matters. At first sight, it seems that the formulation of Definition 1.1.4 is different from a control account. In particular, it does not impose any conditions on A's actions or autonomy, except on the definition of privacy invasions - which are privacy diminishments lacking *A's consent*. Thus, in Véliz account, A's exercising control is switched to A's exercising consent. This is supplemented with explicit conditions (i), (ii) and (iii) for B to fulfil in order to not violate A's right to privacy, not only considering actual privacy diminishments but also counterfactual ones. In this sense, Véliz's account is much more refined from a generic control account of privacy, which tend to simply state that for A to have privacy, A must have control, without imposing explicit conditions on external people for ensuring that control. Finally, note that Véliz explicitly accepts that the right to privacy is not necessarily the most important right to respect in every scenario. Instead, there are cases in which privacy will be invaded justifiably, such as in light of a different prevailing right or value. In that case, someone does not violate but rather *infringes* the person's right to robust privacy. For example, if one has reasonable suspicion and evidence that there is a case of domestic violence in a house, then the right to privacy of the couple must be infringed so that social services can help the victim.

In this section, I have discussed and defined personal matters and the distinction between descriptive privacy (associated with privacy losses) and the normative right to privacy (associated with failures to respect the right, violations, and infringements of the right). In the following section, I will survey and discuss the main theories defining privacy in the literature to further defend the definitions presented in this section.

1.2 Main accounts of privacy

The literature on privacy is vast and multidisciplinary. Research on privacy ranges from developments in encryption (Black 2002) and other techniques like differential privacy in theoretical computer science (Dwork 2006), to investigations on the value of privacy for consumers in economic settings (Varian 2009). Different conceptions and definitions of privacy exist depending on the discipline and the subject of study. I will not attempt to cover every theory of privacy developed in every research discipline. Instead, I will present here the main theories of privacy from the philosophical literature that aim at investigating privacy's nature.

The most common and prominent ways for defining privacy in the philosophical literature are those conceiving privacy in terms of control and in terms of access.[7] One can identify clusters of theories for each kind depending on whether they consider control or access to be the main feature to define privacy. In this section, I will provide an overview of the main proposals for each account and discuss their suitability. I will finish this section by embracing privacy's multi-faceted nature; a full account of privacy necessitates its definition to incorporate both control and access elements, along the lines of the Hybrid Account of Privacy from Véliz (2024). To do this, I will argue and

[7] Although I only discuss control and access definitions of privacy, alternative accounts have been conceptualised outside of this paradigm. Warren and Brandeis (1890)'s account served as an influential first contemporary defence of the right to privacy, which deemed this right to be a special case of "the right to be let alone" (Gavison 1980, p. 437). This condition has been argued as neither sufficient nor necessary for privacy (see Thomson (1975, p. 295), Parent (1983, p. 342) and Véliz (2024, p. 49)). Nissenbaum (2010) posited a theory of contextual integrity, wherein privacy is argued to be "a right to *appropriate* flow of information" (ibid., p. 127), where the appropriateness is given by the context in which the information is taken, processed and, in general, flows. Véliz (2024, pp. 69–70) has critiqued this theory of privacy as being too vague and too permissive with social norms. Lastly, Daniel Solove (2002) identified privacy as a Wittgensteinian family resemblance concept. By considering as hopeless the search for a unique characteristic on every aspect of privacy, Solove argued instead for a theory that connected the different aspects of privacy in a taxonomy (Solove 2006, p. 484). See Véliz (2024, pp. 70–73) for a rejection of Solove's account by arguing that it is unable to make a useful distinction between things that are concerned with privacy and things that are not, showing that it is an unsatisfactory account for defining privacy.

defend that the descriptive nature of privacy is better defined in terms of access and the normative side with a conception similar to a control condition.

1.2.1 Control accounts

Let A and B be distinct individuals and let p be a personal matter of A. Then,

Definition 1.2.1 (**The Control Account of Privacy**). A has privacy with respect to B over p if and only if A exerts some form of control over p with respect to B.

Notice that this definition is intentionally vague and broad since it corresponds to a generic formulation, wherein individual A is required to exercise an *unspecified* form of control. This is deliberate as the goal of this definition is to encapsulate all of the different control accounts in the philosophical literature. Thus, proponents of the Control Account (CA) follow variants of Definition 1.2.1. I will now present some of these proponents, which differ on the exact form of control over p that they deem necessary for privacy protection, and on the nature of p.

An influential account of privacy as a form of control comes from Alan Westin (1967), who discussed the value of privacy considering the increasing prevalence of new, cheaper, and easier-to-use surveillance techniques in North America. Westin defined privacy as "the claim of individuals, groups, or institutions to determine for themselves when, how and to what extent information about them is communicated to others" (ibid., p. 7). In the same spirit, Beate Rössler (2005, p. 71) and Adam Moore (2003, p. 216) have both argued that privacy is the "control over access" to p by external people.[8] Now, while Rössler has the protection of autonomy in mind when defining privacy, Moore focuses on other general reasons to value privacy, such as "human flourishing or well-being" (ibid., p. 223). Thus, these authors have a particular focus in defining privacy with respect to the goods, values, and rights that privacy protects. This is a necessary step for providing a normative theory of the right to privacy. In fact, I will argue in this section that control accounts are especially successful as normative accounts since they have the protection of certain values (in the form of control) as a central point in the definition.

Fried (1968, p. 483) conceptualises privacy as "control over knowledge about oneself", where it becomes relevant whether someone acquires a *general* or a *detailed* fact about us, with the latter being much more intrusive with regards to privacy (ibid., p. 483). There are at least two problems with Fried's theory. First, even though Fried was hoping to clarify that not every acquisition of information about someone will lead to a privacy loss with the distinction between general and detailed facts, this distinction is not actually relevant for a privacy loss. For instance, a transsexual individual might feel that their birth sex is extremely personal and wish to keep it private (in a country discriminatory towards transexual people), even though their sex might be considered a general fact about them. Not all information about us is relevant for privacy, but what is relevant exactly? Instead of differentiating between general and detailed facts, I have chosen to consider to deem the relevant information for privacy as *personal* - not referring to something detailed, but to something that can put us in a vulnerable position. I will further discuss the nature of personal information in Section 2.1. The second issue with Fried's proposal is that he only defines privacy in terms of personal information, restricting his definition to capturing only the informational side of privacy. He conceives B as diminishing A's privacy when B knows about A's *information* without A's control. This conceptualisation is insufficient because personal space is relevant for privacy as well. It is important to note this because, even though this thesis mostly focuses on exploring privacy in terms of personal information, it is crucial to remember that this is only a portion of what the whole concept of privacy should encapsulate. To see this, consider the following. Someone

[8] Note that "control over access" does not entail a dual account combining *both* control and access elements. Instead, this formulation works as a control account where one exercises control over who, how, and under what circumstances another person accesses one's personal matters.

B can invade another person A's privacy simply by B being in A's personal space even if B cannot acquire any information. For instance, imagine that your sibling enters your room as a teenager. Even if your sibling is unable to sense anything, such that they cannot see, hear, smell, taste, or touch anything, it may still feel like the mere presence of them in your room is an invasion of your personal space, and so, a right to privacy violation.[9]

So far, I have presented three of the main CA proposals from the literature, considering privacy to be conditioned to someone exercising control over their information, and over factors such as who accesses it and how. I will now discuss two arguments that favour CA as a general normative account of privacy, and two objections that critique the role of CA as a descriptive account. Interestingly, one of the arguments motivating the CA as a normative account also serves to critique the CA as a descriptive account of privacy via what has been defined as *threatened loss* scenarios. This discussion will show that, if one shares the intuition in this section, control accounts can only respond to privacy's normative side. The first argument in favour of the CA as a normative account, illustrated with the following example, has to do with its ability to capture the moral difference between a situation where someone, A, willingly gives out their personal information p, and a case where p is stolen without A's permission.

Example 1.2.1 (**Credit card**)**.** Consider the following two situations:

(a) A shares their credit card information p with B, their years-long partner, to buy holiday tickets.

(b) C, a stranger to A, hacks A's travel account and steals their credit card information p.

The CA can explain the difference between the two scenarios in Example 1.2.1. In (a), A exercises their autonomy in sharing some of their personal information with a person they at least seem to trust, whilst in (b), A is the victim of identity theft. In both cases, external people have accessed and learned A's credit card information, but the identity theft scenario is much more striking normatively with regards to privacy. For instance, following an instrumentalist view of the value of privacy,[10] one may be worried about A's privacy since situation (b) is likely to lead to fraud and economic loss. A crucial difference between the two situations for CA endorsers is that in (a), A controls the personal information that is accessed, who accesses it and how. In contrast, A does not control any aspect of the situation in which their personal information is acquired in (b). Thus, in general, a CA proponent can explain that A has their privacy invaded in (b) and not in (a). Notice that this intuition resembles that from Example 1.1.2, where we sense an important distinction between the moral implications between the actions of someone eavesdropping a conversation and the friend who is part of that conversation. In summary, if one shares the intuition that a CA can account for the moral distinction between (a) and (b) in Example 1.2.1 and situations (1) and (2) in Example 1.1.2, then the CA is argued to be capable of acting as a normative theory for privacy.

Consider a second argument in favour of the CA that narrates a so-called *threatened loss* scenario.

Example 1.2.2 (**Threatened loss via Alexa**)**.** B hacks into A's Alexa speaker without A's awareness. A's Alexa speaker records A uninterruptedly. B stores all of A's voice recordings f_1,...,f_n in their computer. B never accesses f_1,...,f_n.

This is an example of a threatened loss scenario. It is called this way because privacy is threatened as A does not control the storage of f_1,...,f_n, but a privacy loss does not seem to be *fully* realised

[9] Thanks to Elizabeth Ventham, who presented this example in her talk at the online *Philosophy of Privacy* Workshop, organised by the University of Salzburg on the 16th and 17th of July 2024.

[10] An instrumentalist view of privacy is one where someone considers that privacy is valuable because it protects other goods that we tend to consider to be valuable.

since B never *accesses* the files.[11] A's normative right to privacy is violated, and control accounts offer a plausible reason - it is due to A losing control over their information. This renders CAs as plausible contestants to portraying the right to privacy, but can they explain why it feels weird to say A lost privacy if B never accesses the files? Here, the distinction between descriptive privacy and normative privacy becomes relevant. In particular, I would argue that control accounts do not manage to *completely* explain our intuitions from threatened loss scenarios. In particular, threatened loss scenarios portray cases in which someone A's descriptive privacy is threatened but not lost (since no one actually learns about A's files). Still, quite intuitively, this situation violates A's normative right to privacy - whether one explains it because A loses control over their information, or because A's descriptive privacy is threatened by B and it becomes very easy to diminish it. This is in vein with the analysis from Véliz (2024, pp. 144–147) of these threatened loss scenarios. It follows, then, that even if CAs are plausible contestants for describing the normative right to privacy, threatened loss scenarios present an objection for considering the CA as a plausible descriptive account of privacy. In particular, control accounts are too broad to be descriptive theories of privacy since they predict privacy losses in cases where they have not occurred (Parent 1983, p. 345).

One can also critique the CA as a descriptive account of privacy by showing it is too narrow since there are cases in which, even if someone exercises control, they still lose privacy. Consider the following example showcasing this problematic.

Example 1.2.3 (**Voluntary disclosure on social media**)**.** A is a user on social media. Regularly, A willingly discloses audiovisual content which includes some of their personal information p and personal space s. B is a faithful follower of A. B watches every piece of audiovisual content shared by A.

A control account proponent might argue that A does not lose privacy in this example because they are in complete control over how they share and portray p and s through their social media account. With this line of argument, Andrei Marmor (2023) has argued that social media provides us with a platform where users get to enjoy excessive privacy. Marmor differentiates between the notion of *data protection*, which he takes to be a protection of property rights, and the notion of *privacy*. He argues that the main interest protected by (the right to) privacy is that of "having a reasonable measure of control over ways in which we present ourselves to others." (Marmor 2015, p. 15). Since social media allows us to easily conceal certain parts of ourselves and reveal others, Marmor considers it especially "conducive to privacy" (Marmor 2023, p. 576), though in expense of "authenticity". This argument responds to the intuition that, in social media, people can refine their life story with the use of photo-editing techniques, and by magnifying their life accomplishments and hiding their failures. However, although social media may seem to give us more control over how we present ourselves, it is easy to see that control over self-presentation is not sufficient for privacy.

Lanzing (2016) argues that the kind of exposure that social media facilitates is detriment to our privacy, as it provides an excessive exposure to us. In particular, we cannot foresee who may gain access to the photos we share on the internet once we upload them. This in turn eventually undermines our autonomy. Lanzing's argument focuses on the excessive transparency provided by apps that monitor our daily steps, our daily calories, the time spent on our phone, and so on. These *quantified information* may seem to provide us with more autonomy, in the sense that we can make more informed decisions about how to change our lifestyle. However, the information collected in these apps can be accessed by many companies, and shared with other third parties.

[11] For the present purposes, for B to access the files, it means for B to click on the files, play them and form some form of justified true belief about A depending on the content of the files. Since B never plays the files by assumption, the access is not realised.

This has the potential to undermine our autonomy in much stronger ways, such that the gain in autonomy from using these apps to *change our lifestyle* becomes negligible in comparison. As we can see, this shows that the CA can be used as a normative theory because of its connection to autonomy. However, it is still insufficient from a descriptive point of view. To see this, imagine that social media user A were to actually choose to share their credit card information because they want to portray themselves as a wealthy and reckless person. It seems uncontroversial to say that A loses privacy in this case. In this sense, CA accounts cannot explain that people can lose privacy if they act recklessly and over-share, even if they are technically in control. In particular, although A may control their self-presentation when sharing the amount of money of their back account, and the luxurious street in which they live in, A is losing privacy when people access their personal matters because they are putting themselves in a vulnerable position where they could be robbed for instance.

In the discussion of the CA so far, I have attempted to show its benefits as a normative account of privacy and its deficiencies as a descriptive theory. Let us now turn to discussing access accounts of privacy.

1.2.2 Access accounts

Let A and B be distinct individuals and let p be A's personal matters. Then,

Definition 1.2.2 (**The Access Account of Privacy**). A has privacy with respect to B over p if and only if B does not have (some form of) access to p.

Notice that, within an Access Account (AA), the fact that B "does not have (some form of) access to p" may mean either that B has *limited access* to p, that p is *inaccessible* to B, or that p simply *non-accessed* by B. Let us discuss this options in turn. Gavison (1980) identifies the definition of privacy with the first interpretation of the AA. Gavison (ibid., p. 428) conceives privacy as a "limitation of others' access to an individual". Gavison takes privacy to be in limiting access of something to other people because she posits that speaking of privacy in absolute terms is nonsensical. This is coherent with earlier claims since living a fulfilling life in society tends to mean that we enjoy a certain degree of privacy and must sacrifice another, whether it is for building relationships, taking care of routine administrative tasks, or seeking help from others. In that sense, Gavinson inferred that to define privacy as complete lack of access to others must also be nonsensical. However, even if Gavison's rejection of privacy as absolute is correct, her formulation of privacy as limited access does not yield desirable results. Parent (1983) argues that limited access theories do not manage to cover all cases in which privacy is diminished, that is, they are too narrow. To see this, consider the following example.

Example 1.2.4 (**Bad signal**). B taps into A's phone and overhears their conversations. However, there is bad signal in the network during the tapping so B can only hear every second or third word from A's conversation. B is still able to infer what is being said from the conversations and, in particular, they infer personal information p.

Although B has limited access to A since their eavesdropping is hindered by a bad signal, B accessed p by inference and, therefore, A's privacy is still diminished. This counter-example shows that limiting access to something does not ensure privacy, rendering it an insufficient condition. Someone could reply by noting that what is actually relevant for whether A's privacy is diminished over p is that B's access is limited in such a way that B is incapable of accessing p. This conceptualisation appears from the second interpretation of the AA - A has privacy over p with respect to B if p is inaccessible to B. Reiman (1995) is a proponent of such interpretation as he conceives privacy as "*the condition in which others are deprived of access to* [A]" (ibid., 30 emphasis in the original).

However, there is a problem with requiring that a personal matter p is inaccessible for privacy to be protected since this conception collapses into a control account of privacy. Requiring that p is inaccessible means that there is a *way* to ensure p is non-accessed for some period of time. Otherwise, if there is no way of ensuring this, we simply could not claim that someone cannot access p. That is, for p to be inaccessible we require some form of control over p and its access. Otherwise, we lose control over its access and, with that condition breached, following Reiman, our privacy. Véliz (2024, p. 76) similarly notes that "when something becomes accessible to others we lose control over it". Therefore, rendering inaccessibility as a condition for descriptive privacy is similar to positing a control account. For this reason, Mainz and Uhrenfeldt (2020, p. 292) argued that proponents of the access account of privacy must not require for people to be unable to access a persona matter for it to be private, but instead privacy should be defined in relation to an object being non-accessed, without making claims about its future.

Finally, we are left with the third interpretation of the AA account, which posits that A has privacy over p with respect to B if p is *non-accessed*. This is the interpretation that I will endorse throughout this thesis. A proponent of this interpretation is Parent. Parent (1983, p. 346) conceives privacy as "the condition of a person's not having undocumented personal information about himself known by others". We can see that Parent is a proponent of privacy being ensured by non-access where he assumes that *accessing* p is *acquiring knowledge* about p. We will see in the next chapter that this assumption is non-trivial, and we will carefully inspect what it means to access p in light of privacy. Now, before we go into assessing the plausibility of the AA as a descriptive theory of privacy, given this interpretation, I will make an important note that arises from Parent's formulation.

Parent's theory relies on the difference between *documented* and *undocumented* information. In Parent's conception of privacy, only undocumented information is worthy of protection, where undocumented information is that which is not present in the public domain, such as public records, newspapers, and, in extension to present times, the internet. The problem with this distinction has to do with confusing what privacy is with the distinction between the public and the private sphere. The issue with this confusion is illustrated with the following example by Moore (2010, p. 19).

Example 1.2.5 (**Walk in the park**)**.** Imagine A is walking in the park. In doing so, A is exposing all types of information about their appearance for bystanders to see, from their hair colour to their approximate height and weight. By walking, A inevitably leaves some of their biological matter as they walk by, like hair strings, sweat, and skin.

In this example, A's hair strings become part of the public sphere as A walks. Moore (ibid.) notes that, if one were to follow Parent's definition that privacy only protects undocumented information, then the genetic information that one could find out from A's hair strings is not protected by privacy. By relying on the distinction between documented and undocumented information, Parent confuses what is public information with that which is in the public sphere. My body is in a public space when I leave my house but that does not mean that this render it okay for anyone to access. Consider another example that shows that what should be public and what should be private is not connected to what occurs in the private sphere, e.g., one's home, as opposed to what happens in the public sphere, e.g., a public street or beach.

Example 1.2.6 (**Injury**)**.** A suffers a severe injury in a heavily transited street. An ambulance arrives and the paramedics need to force their shirt open to apply necessary medical procedures. Bystander B records A's naked torso throughout the entire procedure and publishes it on the internet.[12]

[12]Example from (Véliz 2024, p. 36).

If one follows Parent distinction of documented and undocumented information, B does not diminish A's privacy because A is in a public place, and A cannot have privacy regarding things that are in the public eye. I think that its is fairly intuitive that this is wrong. Privacy must not be equated with the private sphere, and neither should the public be equated with the public sphere. This is most strikingly grounded on pragmatic considerations such as the fact that we should be able to enjoy a private conversation with our partner in a restaurant without eavesdroppers, to do sports in the park without being recorded, and to walk down the streets without our genetic or location data being retrieved by external parties.

So far, I have argued that the plausible AA of privacy should consider privacy as a condition that person A has when their personal matters p are non-accessed by B (as opposed to p being inaccessible or B having limited access to p). Additionally, p is not only protected if it is undocumented information, but it is protected as long as it is personal information (or, to be precise, personal matters). Now, Véliz (2024, p. 79) has recently proposed an access account of privacy that agrees with these intuitions. My own formulation of her definition is the following, which I also presented in the previous section. I will now defend it as a suitable descriptive notion of privacy. Let A and B be distinct individuals, and let p be a personal matter of A.

Definition 1.2.3 (**Descriptive Privacy**)**.** A has *descriptive privacy* with respect to B over p if and only if B lacks access to p.

The main advantages of the AA relate to its ability to deal with the objections from the previous section that critiqued CA's unsuitability as a descriptive theory of privacy. Let us begin by noting that the AA can account for the *threatened loss* scenarios from the previous section. Recall that these scenarios capture cases where someone's personal matters p are at risk but, for some reason, they have not been actually heard, read, or learned about. In Example 1.2.2, B hacks A's Alexa speaker and collects all of A's recordings; however, B does not actually hear these recordings. Within an access account, privacy is not lost in a threatened loss scenario because A's personal matters are simply *not* accessed.

Additionally, the AA can explain the privacy diminishment that people suffer in cases of *voluntary disclosures* of personal information from Example 1.2.3. Therefore, if someone A enjoys posting a lot about their personal life on social media, and, one day, A posts their credit card information, then A lose privacy when their credit card information is accessed by another person, even if A could be said to be in control of their information. Thus, the AA account provides a succinct and successful theory of the descriptive theory of privacy.

Now, one can also critique the AA as being insufficient for capturing the normative side of privacy. In contrast to the CA, it cannot account for a difference between two people B and C acquiring the same personal information through different means as illustrated in Example 1.2.1. While B learns about A's credit card information because A willingly tells them, C steals A's credit card information. An access account cannot capture the *manner* in which a personal matter is accessed, it only cares about the fact that the access itself occurred. Additionally, the AA cannot account for the normative wrong which occurs when B compiles and collects A's recording files from their Alexa in Example 1.2.2. B has purposefully created a situation in which it would be very easy to diminish A's privacy. However, AA does not capture this difference, and fails to explain the intuition that A's normative right to privacy was violated by B. As a last objection to the CA, let us consider Fried (1968)'s famous *desert island* example, where he highlighted a counter-intuitive consequence arising from AAs. Consider the following modified example from Fried (ibid., p. 482).

Example 1.2.7 (**Desert island**)**.** A is stranded in a desert island. No one can access any of A's personal matters. For this reason, according to the AA, A enjoys full and perfect privacy.

Fried's purpose with this example is to display the ironic consequence of access accounts of privacy, where having full and perfect privacy is simply undesirable. Indeed, speaking of *A enjoying* privacy when stranded in a desert island is ridiculous. However, it is not an actual counterexample to the AA because the AA does not posit any moral claims on privacy, i.e., it does not account for whether having or losing privacy is good or bad - it simply gives an account to what it means to have or lose privacy. This is related to the AA being suitable as a descriptive account. So, indeed, Example 1.2.7 portrays a strange consequence of a theory of privacy based on the AA, but it does not show its unsuitability as a descriptive theory (Reiman 1995, p. 30; Parent 1983, p. 349).

In this section, I have presented and discussed the two main accounts of privacy in the literature, control and access accounts, while I weighted their suitability in attaining to our intuitions with regards to the descriptive and the normative conceptions of privacy. I attempted to show that control accounts are better suited to define the normative right to privacy, as they are motivated by moral considerations regarding privacy, while access accounts manage to capture privacy's descriptive and more neutral side, managing to define privacy before one considers if privacy in itself is good or bad. Let us further motivate the distinction between descriptive and normative privacy in the last section of this chapter.

1.3 The descriptive and normative sides of privacy

Throughout this chapter, I have motivated, defined, and discussed privacy as a two-fold notion, with a descriptive and a normative side. I will now finish this chapter by further arguing for the necessity of having two notions of privacy that are differentiated, where each of the notions is necessary to capture privacy's full meaning. To begin defending the differentiation between descriptive and normative privacy, I will present two arguments from Skopek (2020, pp. 8–13) that show that distinguishing between descriptive privacy (related to privacy losses) and normative privacy (related to privacy violations) helps us avoid falling into at least two kinds of misdirected philosophical discussions.

First, Skopek argues that philosophers sometimes give *mistargeted critiques* to certain accounts of privacy by failing to see that they are solely concerned with the descriptive notion of privacy, or solely concerned with the normative side. This is perfectly coherent with the entire discussion in the previous section, where I have argued for the deficiency of control accounts as descriptive theories of privacy, and access accounts as normative ones. So, to entertain directed philosophical discussions, one should settle whether their theory is descriptive or normative and, from there, another person can support or critique such theory within this angle. Otherwise, one may find themselves critiquing an access account of privacy (that aims to be a descriptive account of privacy) because it fails to account for our normative intuitions - stagnating the discussion due to its triviality.

Second, Skopek argues that authors have also shown a *misguided scepticism* about providing a unified account of privacy because of failing to differentiate between the viable unification of descriptive privacy and normative privacy. Skopek (ibid.) notes that the accounts of privacy from Thomson (1975) and Solove (2002) are two examples of this kind of scepticism. Let me focus on Thomson (1975), who famously argues for a *reductionism* of the right to privacy. Thomson's reductionism arises from considering the right to privacy as, first, having no use in itself and, second, being wholly derivable from other rights, such as the rights to property or the right against physical intrusions. Now, Thomson's analysis is important to raise a discomfort with the concept of privacy, and it is still being fairly discussed in the privacy literature. However, it has been argued to fail as an account to capture the nature of privacy and the right to privacy. In particular, reductionism has the undesirable effect of entailing that we have rights for things that we would not consider normally to receive such moral standing (Véliz 2024). Since Thomson claims that

the right to privacy is derivable from other rights, there must be rights that protect all kinds of things, such as a right against unfounded gossip. A proliferation of meaningless rights seems to be undesirable consequence of Thomson's theory.

Now, by differentiating between descriptive and normative privacy, we can see that Thomson's reductionism is misguided. On the one hand, there is indeed a plurality within the right to privacy that cannot be reduced, as there are various kinds of things (personal matters) that are protected from other people's access due to a wide variety reasons (such as the preservation of integrity, happiness, calmness). However, unlike Thomson, it seems reasonable to me that the plurality of the right to privacy is actually desirable exactly because it allows us to account for different moral intuitions about what privacy should protect. On the other hand, even with the plurality of the right to privacy, we can find theoretical coherence in another notion related to privacy: descriptive privacy. Skopek (2020, p. 13) argues that descriptive privacy allows us to contain all actions that relate to privacy losses *outside* of our moral intuitions. Descriptive privacy unifies the notion of privacy as it does not attempt to say whether privacy is good or bad in certain situations and circumstances. Instead, descriptive privacy provides a way to say that a person has lost privacy, as someone has accessed their personal matters (within an access account), and this can be judged as good or bad outside of the paradigm with the normative right to privacy. This way, the notion of privacy can enjoy theoretical coherence from its descriptive side (ibid.). In summary, I just presented two reasons from Skopek (ibid.) for which one should differentiate between descriptive and normative privacy, i.e., to avoid philosophical discussions that focus on *mistargeted critiques* and *misguided scepticism*.

As a side note, it may be important to mention that, although I have throughout stressed the importance in distinguishing between descriptive and normative privacy, it is also important to note that this distinction is a *somewhat* illusory one. Descriptive privacy will always need to attend to our values and intuitions, so that it is grounded in practical reality. Therefore, descriptive privacy is not (and should not be) completely normatively neutral. As I understand it, descriptive and normative privacy are differentiated in the following way. On the one hand, descriptive privacy captures the kind of actions that have the potential to become right to privacy violations. In other words, descriptive privacy captures the actions that, as a culture or society, we have identified as placing us in vulnerable positions (this is what guides which matters - information or space - are personal). In her book *Privacy is Power*, Carissa Véliz (2020) argues that we need privacy because it protects us from power abuses. We could be wronged by someone using our personal information or by them invading our personal space. However, it is not the case that every time that we are in this vulnerable position, people take advantage of us. Descriptive privacy is an attempt to abstract away from the ways in which someone wrongs us by abusing the power that knowing certain things about us gives them. Instead, within descriptive privacy, one identifies the actions by which this power abuse *could* occur. On the other hand, normative privacy identifies wrongs arising from attempts or actual diminishments of our descriptive privacy. These wrongs can take place before our information or space is even accessed, when they are accessed, or after they are accessed by, e.g., improper use. Given these intuitions, let us present descriptive and normative privacy for a final time in this chapter.

Descriptive privacy

A *descriptive account* of privacy must explain the phenomena of privacy diminishments (and losses). In its descriptive conception, privacy is a state or condition (Moore 2010, p. 14), which one can have, one can lack, or one can lose. In some sense, descriptive privacy must capture the actions that make us vulnerable regarding things *about* us. For example, telling our secrets, letting someone be physically close to us, or sharing our health-related information. The tricky thing is that

what makes us vulnerable changes as society changes. People from the LGBTIQ+ collective hide in dictatorial regimes because their sexual orientation can put their life in danger. This is why LGBTIQ+ Pride parades are so important in inclusive liberal democracies. These parades are a way to show that people should not be scared of sharing their sexual orientation as discrimination for sexual orientation should not be a thing. In a utopian non-discriminatory and inclusive society, sexual orientation ceases to be a piece of personal information because this fact would no longer put people in a vulnerable position.

I have argued that a descriptive theory is best described within an access account of privacy. As I presented in Section 1.1, A has privacy over their personal matters p with respect to external person B if B lacks access to p. This definition allows us to account for cases in which privacy losses take place even though we have not been wronged. This is illustrated in Example 1.1.1 and Example 1.2.3. Two friends building a relationship and sharing personal details with each other in confidence is not a violation of privacy, even though their descriptive privacy is indeed diminished. Similarly, if someone A voluntary discloses their personal information, A is not wronged (their right to privacy is not violated) even though A is diminishing their own descriptive privacy.

The normative right to privacy

A theory of *normative* privacy should account for the value of privacy in situations that are morally sensitive. In this sense, the normative conception of privacy is usually described in terms of a *right* to privacy. As argued in Section 1.2, a control account of privacy is better suited for these normative endeavours than an access account. However, one may recall that I presented Véliz's right to robust privacy as my endorsed normative theory of privacy in Section 1.1. This is because control is not the best suited condition to ensure that the right to privacy is respected. To see this, one needs to understand that control is *not* necessary for respecting the right to privacy because one can lose control without their right to privacy being violated. This fact is exemplified in a certain version of threatened loss scenarios presented in Section 1.1.

Recall that a threatened loss occurs when person A has some of their personal information p taken by person B without B accessing p. In Example 1.2.2, A has been wronged by B because B robbed A's recordings p from A's Alexa. Now, let us consider a threatened loss example from Macnish (2018), where A's right to privacy is not violated to show that A can lose control over p with their right to privacy respected. Imagine that A forgets their diary in a café. An hour later, A realises they left their diary there, so A returns to the café. When they enter, they find their diary on the table of stranger B. B, after seeing A, provides them with their diary after assuring that they did not read it and instead kept it safe on their table because they are aware of the importance of a personal diary. Let us assume B is speaking the truth. Now, in the hour that A is not in possession of their diary, they can be said to have lost control. However, A's right to privacy has not been infringed, disrespected, or violated, since B did not read the diary. In this case, losing control has not entailed a failure to respect the right to privacy. Therefore, control is a condition that is too broad for the right to privacy. In Chapter 3, I will present, explain, and critique Véliz's right to robust privacy as a plausible contendant for capturing the normative right to privacy.

1.4 Conclusion

In this chapter, I have began investigating the nature of privacy by motivating and defending that our privacy intuitions respond to two notions that must be differentiated: a descriptive and a normative notion of privacy. To illustrate this, in Section 1.1, I have discussed some preliminary examples and I have presented the definitions of descriptive privacy and the normative right to

privacy based off the Hybrid Account of Privacy from Véliz (2024). Then, in Section 1.2, I have presented and discussed the two main accounts defining privacy in the literature: access accounts and control accounts. I have argued that control accounts are more suitable for normative considerations, while access accounts adjust better to our descriptive intuitions. Finally, in Section 1.3, I have argued that *control* is not necessary for a normative notion of privacy, and I have further discussed the necessity for differentiating between descriptive and normative privacy, and elucidated their nature.

Therefore, the main goal of this chapter has been to clarify the nature of privacy by explicitly defining and presenting a distinction between descriptive privacy and the normative right to privacy. A summary of the main characteristics of these notions is presented in the following table.

Account	**Descriptive privacy**	**Normative right to robust privacy**
Definition	1.1.2	1.1.4
Related Actions	Privacy losses/privacy diminishments (1.1.3)	Infringements of the right (3.1.3) Failures to respect the right (3.1.1) Violations of the right (3.1.2) Invasions of privacy (privacy losses without consent)
Concerned with	Actual world	Actual and relevant possible (counterfactual) worlds
Assessment with respect to	Personal matters (1.1.1) (*path-independent*)	Manners of getting to personal matters (*path-based*)
Discussed in	Chapter 2	Chapter 3

Table 1.4.1. Summary of the definitions, characteristics, and actions related to descriptive privacy and the normative right to privacy.

In the following chapters, I will continue elucidating the nature of descriptive and normative privacy, respectively, by focusing on their connection to epistemology.

Chapter 2

The Epistemology of (Descriptive) Privacy

In this chapter, I will analyse descriptive privacy through the lens of epistemology. In particular, my main goal is to analyse and model the phenomenon of accessing a piece of personal information using epistemological notions. I will refer to such endeavour as 'the epistemology of privacy'. I will focus on the analysis of informational privacy, i.e., privacy with respect to personal information, and leave the complications of personal space aside for two reasons. First, for practical purposes, focusing on the dimension of personal information, and excluding personal space, reduces the scope of the analysis of this chapter and allows me to take a deeper focus on the epistemological issues arising from personal information. Second, it is more straightforward to link personal information to the acquisition of knowledge or some other epistemic state, than it is to elucidate the implications, epistemic or otherwise, from invading someone's personal space. Thus, to avoid getting lost into the particular phenomenological or epistemic phenomena arising from personal space invasions, I will restrict to personal information. In this chapter, whenever I write *privacy* I refer to *descriptive informational privacy*.

One may wonder whether analysing privacy within epistemology makes sense at all. Recall the definition of descriptive privacy provided in the previous chapter. According to this definition, individual A has descriptive privacy about personal matter p with respect to individual B if and only if B *lacks* access to p. Restricting ourselves to personal information, this seems to imply that protecting privacy entails a *restriction* on knowledge or, more generally, on epistemic access. In fact, Goldman (1999, p. 173) argues that privacy is not an interesting notion to study within epistemology because privacy is solely concerned with "knowledge curtailment", while epistemology should study acts that enhance knowledge. Later on, Goldman (2002, pp. 218–220) conceded that epistemology could broaden and gain insights as a field from studying notions that restrain knowledge.

Fallis (2013) disagrees with Goldman. Fallis thinks that privacy can benefit from being studied as an epistemological notion. However, unlike Goldman, Fallis argues that this is due to privacy being also able to lead to an increment in knowledge acquisition, not only to knowledge curtailment. To illustrate this, Fallis discusses the case of library patrons (ibid., p. 154). Libraries in the US are committed to the privacy of their patrons by protecting their borrowing history, which is only released in the presence of a subpoena (Garoogian 1991, p. 230). Additionally, in the advent of information technologies, librarians adjust the default options of the library's computers to ensure that the browsing history is deleted after use (Lamdan 2013, p. 134). This way, libraries ensure that patrons are free to read and disseminate information without the fear of surveillance or interference from the government or other organisations (ibid., p. 131). Of course, one could note

that this actually leads to the knowledge curtailment of FBI agents that use surveillance practices to monitor terrorism (Fallis 2013, p. 154). However, the effectiveness of mass surveillance in the fight against terrorism is controversial and contested (Véliz 2024, p. 147), and surveillance can cause chilling effects on users, i.e., it can lead to users self-censoring in the information they access as they worry that external parties can monitor what they search, the books they borrow, or their ideas (Garoogian 1991, p. 229). Thus, the absence of privacy in the form of surveillance can lead to user's self-censoring and abstaining from searching particular topics by fear of repercussions. In the context of an autocratic government, a lack of privacy is especially dangerous. The censorship arising from chilling effects can lead to an inability to entertain ideas outside the *status quo*, and, in this way, the kind of knowledge one can obtain is highly restricted. In the contrary, the presence of privacy in libraries can avoid this chilling effect, which would likely lead to the acquisition of different forms of knowledge.

One may note that the library patrons' example only shows that a lack of privacy can impede knowledge acquisition and perhaps that privacy can, *in some cases*, be conducive of knowledge enhancement. However, it does not show that privacy always leads to knowledge enhancement, nor than privacy is constitutive of knowledge enhancement. Ultimately, these particularities are not important to motivate the epistemology of privacy. The library patrons example aimed to show that privacy is important for the epistemic freedom of individuals. However, regardless of whether epistemologists find the study of privacy interesting or illuminating for the analysis of knowledge, my aim is to show that privacy studies can gain significant clarity from *using* epistemological notions and investigations. In this thesis, I will pursue this goal in two parts. First, this chapter will elucidate the theoretical clarity that the notion of descriptive privacy can gain with epistemological notions, which is where most theorists have focused their theoretical efforts. Second, in the next chapter, I will propose novel ways in which epistemological notions, such as *knowability*, can aid the analysis of the *normative* dimension of the right to privacy.

This chapter is organised as follows. In Section 2.1, I will present and expand on the epistemological analysis of privacy by Blaauw (2013), where he describes privacy as a relational state between a subject, a set of personal propositions, and a set of individuals. This discussion will allow us to discuss the gradable nature of privacy. Then, in Section 2.2, I will inspect the nature of privacy losses, in order to elucidate the actions that must not occur for someone's privacy to be preserved. For this purpose, I will inquire on the conditions necessary for epistemic access. First, I will take the acquisition of knowledge as a simple example and inspect whether weaker epistemic states than knowledge can also lead to a privacy loss. For this, I will inspect the necessity of the conditions of *truth*, *belief*, and *justification* (and other forms of *epistemic merit*) in Sections 2.3, 2.4, and 2.5, respectively. I will survey the main accounts from the epistemology of privacy literature throughout, and I will provide the first argument against Véliz (2024)'s novel theory, where she posits that unjustified true beliefs can be sufficient for a privacy loss. Instead, following Skopek (2020), I will argue that epistemic access occurs only if an agent acquires a true belief about someone else's personal information, where there is a link between truth and the belief that satisfies a form of *epistemic merit*, e.g., justification, reliability, causality.

2.1 The relational state of privacy

Blaauw (2013) describes the state of having privacy as a ternary relational state between (1) a subject A, (2) a set of individuals I, and (3) a set of personal propositions P. P is composed of personal propositions $p_1, ..., p_n$. Note that, whenever I define a notion within descriptive privacy, I will refer to personal proposition p as an abstract and undefined amount of personal information. Similarly, I will refer to an undefined single individual B within set I. In the following, let us

explore each of these *relata* in detail.

2.1.1 The subject A and the set of individuals I

The *subject* A is the being whose privacy is under evaluation. The question to consider here is whether a being needs to fulfil a certain criterion to be eligible to have or to lose privacy, e.g., does it need to be alive, or have consciousness? It seems clear that it is meaningful to speak of an adult human being as having or losing privacy. However, can infants have privacy? What about animals, or the deceased (Blaauw 2013, p. 168)? It seems much more likely for people to agree that infants can have privacy than it is for someone to think than the deceased can. When it comes to animals, behavioural and social studies tentatively show that the instincts for privacy in humans have their roots in animal behaviour (Westin 1967, p. 56; Pepper 2020, p. 628; Véliz 2024, pp. 9–15). Animal activists have used this to argue that animals have a right to privacy that humans violate when, for example, gorillas and other animals are visibly distressed when exhibited in zoos (Pepper 2020, p. 629). Upon consideration of the privacy of the deceased, a debate has emerged as a consequence of the current digital age, where it is almost inevitable to have information about oneself on the internet if one is to play a role in society and the workforce. As we use applications, websites, and services online, digital footprints of each of us are created and it is extremely hard—or impossible—to delete most of our data residing in both public and private servers. In this context, the question of whether us, or our loved ones, have any say on what happens to that data after we die has become especially relevant (Allen and Rothman forthcoming). As an example, it is now possible to use generative artificial intelligence tools trained on someone's texts, images, and other information to reanimate celebrities and family members in the form of chatbots or synthetic images and videos. There are many moral concerns surrounding this trend, and one of them has to do with privacy and consent.[1]

My take on this debate is that the question of whether individuals of a particular group can have a right to privacy can only be answered in the context of a normative theory on the value of privacy. For instance, in the context of the right to privacy of animals, Pepper (2020) considers privacy to be valuable for establishing meaningful relationships and to control our mode of presentation. Within this theory of the value of privacy, Pepper argues that animals not *only* have a right to privacy when they are aware that they are being observed, e.g., in the zoo—which is what would be argued by someone who holds that the value of privacy resides in its ability to avoid harm and discomfort caused from surveillance. Instead, Pepper argues that animals have a right to privacy even when they are not aware that they are observed because their ability to form relationships and their ability to control their mode of presentation is violated by covert surveillance. Now, many authors have found the value of privacy to reside in its necessity for people's well-being, creativity, concentration, autonomy, and democracy, among other goods (Fried 1968; Gavison 1980; Moore 2003; Rössler 2005; Véliz 2024). From this discussion, I posit that any being who can be said to be sensitive to these values, is also able to have or to lose some degree of privacy.

One can notice that the normative side of privacy permeates into its descriptive side. Even though I am defining features from the relational state of descriptive privacy, I still need normative intuitions. Recall that privacy is a heavily normative term, and our intuitions about its meaningfulness attend to normative standards. To capture the difference in meaning between saying that a human has privacy, as opposed to a plant or a rock, normative standards are inescapable. As I argued in the previous chapter, this is not a flaw but a feature of privacy. Descriptive privacy is

[1] See (Allen and Rothman forthcoming) for both a survey of the current legal environment on a right to privacy for the deceased and a theoretical stance of how it should look like. Allen and Rothman identify three interests for such right: the interest that humans have in the (1) handling of the image and data of their future deceased selves, and in the (2) handling of the data of their deceased loved ones and family members, and (3) a societal interest to respect the dead.

the notion (related to the normative side of privacy) that remains when, first, we identify the actions that have the potential to wrong us regarding our privacy since they place us in a vulnerable position—e.g., when we tell our secrets or when we have someone standing very close to us—and, second, we abstract away and find the *conditions* for this vulnerability to occur (i.e., an individual accesses our personal information or personal space).

I will now discuss the nature of the set of *individuals* I. These are the individuals (with respect to subject A and A's personal information p) that access p, attempt to access p, or are in a position in which they could easily access p. This definition responds to the needs of the right to robust privacy, as we attend to an individual who not only accesses but *could* access p. In contrast, Blaauw (2013) defines I as the group of people that A does not wish that they access a certain subset of A's personal information P. This definition misses something that I consider crucial of descriptive privacy. In particular, descriptive privacy should capture the intuition that one loses privacy if one confesses a personal secret to someone, even when they are a close friend. As one puts themselves in a vulnerable position that could be potentially exploited, privacy is lost, regardless of whether this vulnerable position is actually exploited or not. Therefore, anyone who is in the position to diminish your privacy can be part of set I.

Regardless, within descriptive privacy, one may also wish to capture the intuition that it is not the same for your sibling or a stranger to access some of your personal information p. In particular, one may think that someone diminishes more, or less of your privacy depending on who it is that accesses p. For instance, if I have two friends G and H, and I know that G loves gossip, while H is very discrete and values privacy, I may feel that I lose more privacy when I tell G about my family problems, than when I tell H. We can capture this intuition by realising that privacy is a gradual notion. As I will argue in Subsection 2.1.3, the relationship between a person and subject A is important for deeming how much descriptive privacy A loses. This responds to the intuition that whoever this person is plays a role on how comfortable A is telling them and how likely it is that may go and tell someone about it. Similarly to the difference that I argued between friends G and H above. Even with this gradual nature, it is important to stress that A always loses privacy to someone if that person accesses their personal information regardless of who they are. However, how much privacy A loses may depend on who this person is.

2.1.2 Personal information P

Let the set of A's *personal information* P be composed of propositions $p_1, ..., p_n$.[2] As a first step towards a definition, recall that the European Union's General Data Protection Regulation defines personal information as "any information relating to an identified or identifiable natural person" (European Parliament and Council of the European Union 2016, May 4, Art.4). Thus, loosely speaking, $p_1, ..., p_n$ must be about A. However, a considerable amount of information that could be said to be about A, or relating to A, is not usually conceived as personal. For example, one would not normally regard the number of strings of one's hair as personal information. Nevertheless, I will defend now that every piece of information about someone is *potentially* personal information, i.e., in the right context with the right amount of information and effort.[3] To begin with a simple example, the number of hair strings can constitute personal information for someone who is undergoing a cancer treatment and who wants to conceal that they are bold as a consequence of chemotherapy with a headscarf or a wig (Kappel 2013, p. 181). In this context, the number of

[2]This is an idealisation to ease the abstraction of personal information. I do not mean to suggest that personal information can *actually* be explicitly detailed and enumerated. Each p may be conceptualised to different degrees of generality.

[3]Matheson (2007) and Fallis (2013) have proposed that there might be propositions about someone that can never be personal, i.e., it would be meaningless to talk about someone having or losing privacy about them. The example that Matheson gives is the fact that someone is self-identical. This is not completely surprising, though, since tautologies are generally considered uninteresting from an information theoretic point of view.

hair strings in someone's hair is personal information. Thus, whether some piece of information is personal information depends on context (Kappel 2013, p. 181).

Many people may be shocked by my claim that any information can become personal in the right circumstances. Surely there are things I do not care that people know about me, such as the colour of my hair, my approximate height, the amount of words I write in a day, the concerts I went to last March, the amount of calories that I ingested in the past week, and so on. I do not necessarily feel that I lose privacy or that I can be harmed in any way if people were to access this information. This is of course partly true, there are all kinds of information about us that we consider trivial and uninformative. However, for any of the information that I listed just now, an external person can judge and update their beliefs according to it. This can eventually lead to agents' finding out personal information through some kind of reasoning or inferential process. For example, imagine that T tells their friend F about their calorie intake yesterday. Correspondingly, F forms belief (A).

(A) T's calorie intake yesterday was 3500 calories.

It so happens that F studied Nutrition and, as T's friend, F also knows basic information about T's body and lifestyle. Thus, F has the background beliefs (B), (C), and (D).

(B) T has a sedentary lifestyle, and is a 30 year-old male.

(C) The recommended daily calorie intake for a sedentary 30 year-old male is around 2500 calories.

(D) If someone's calorie intake is above the recommended one, then they are very likely to develop cardiovascular diseases.

From (B) and (C), F infers the following:

(E) The recommended daily calorie intake for T is around 2500 calories.

Then, comparing (A) with (E) in conjunction with (D), F infers the belief that T is very likely to develop cardiovascular diseases.

Now, imagine that instead of T's friend F, it is a health insurance company that is calculating this inference based on information that T inputs in an app that tracks daily calories intake. Inferences have important privacy implications as they lead to unforeseen privacy invasions.[4] This example aims to show that any information related to us can become personal in the presence of other information—in the example, this refers to F's background information (B) and (C)—alongside with some effort—required to perform the inference that led F to conclude that T is very likely to develop cardiovascular diseases.

Someone could argue that claiming that every piece of information about us can *potentially* be personal is too broad and useless. How do we expect to have any kind of privacy regulation, or be able to obey with moral duties to respect people's right to privacy? Indeed, I agree that it would be ridiculous (and impossible) to try to keep every information relating to us concealed from people constantly. This would also be undesirable for a prosperous society. However, this may just mean that we need to create contextual solutions for assessing when some information about us, although it is not directly personal, should be protected because of the possible implications to finding out more personal information. In fact, in the age of big data, all kind of inferences can be made regarding our *trivial* information. From elucidating our political views to predicting our sexual orientation, this is possible with techniques that aggregate information, spanning from social media to websites visits and event registrations.

[4] Cases where personal data is uncovered from *seemingly* innocuous data often occur in the context of *anonymised* or *pseudo-anonymised* data. A recent example relates to a data breach of the US telecommunications company *AT&T* that encompassed call and text logs from several months (Cox 2024). The metadata from this data breach did not contain names. However, using publicly available tools and lists with people's phone numbers, hackers were able to retrieve call and text records from members of the family of Donald Trump and the family of Kamala Harris (Cox 2025b). This is not only a right to privacy violation, it also presents a national security risk.

I will lastly note that, for the rest of the chapter, I will simplify the concept of a set of personal information P to be personal information p, which can correspond to an indeterminate amount of information. As I will explain in Subsection 2.1.3, privacy is a gradable notion, so speaking of the full set of personal propositions P is unnecessarily exhaustive, given that it is impossible for someone to lose *all* of their personal information with respect to a group of people. Thus, although theoretically nice, these full sets do not stand for anything meaningful in reality.

2.1.3 Privacy is a gradual relation

After exploring each of the relata which are necessary for privacy in detail, it is important to define now how to connect them into one definition of the relational state of privacy. My proposal is the definition of descriptive privacy that I defended in the previous chapter, which provides an account of what amounts to *having privacy*. I have interpreted Blaauw's sets for personal information P and individuals I as encompassing all pieces of personal information about A and all individuals that could diminish A's privacy, respectively. These are not useful notions because, as I will argue here, privacy is gradable. Therefore, I will abstract these notions to be represented by p and B, which denote, on the one hand, specific but indeterminate amounts of personal information about A and, on the other hand, a specific but indeterminate external person that can diminish A's privacy. Now, as usual, let A and B be distinct individuals, and let p be A's personal information. Then,

Definition 2.1.1 (**Descriptive Privacy**)**.** A has privacy about p with respect to B if and only if B has not epistemically accessed p.

In the following sections of this chapter, I will discuss the nature of *epistemic access* via the presentation and discussion of the main views on the epistemology of privacy literature. However, prior to this, I would like to note something important about the concept of privacy. Privacy is not an absolute notion that one simply has or one does not have. It would be impossible for someone to have full privacy if they also enjoy a fulfilling social life, perhaps only a hermit or someone in a desert island could have full privacy.[5] Similarly, it would be impossible for someone to have no privacy at all, as it would mean that every person knew everything about them. People do not have neither the cognitive capabilities nor the interest for this to occur. Thus, in the real world, privacy comes in degrees. Gavison (1980, p. 440) notes that "[i]ndividuals must be in some intermediate state—a balance between privacy and interaction—in order to maintain human relations, develop their capacities and sensibilities, create and grow, and even to survive".

Blaauw (2013, p. 171) notes three conditions that vary the extent of a privacy loss. First, the *number of personal propositions* which A loses privacy over (Kappel 2013, p. 180).[6] It seems reasonable to consider that someone reading your entire diary diminishes more of your privacy than if that person were to read only the first entry. Also, the *number of individuals* that access a certain personal proposition. In fact, it seems that our privacy is more diminished if something is posted on a social media platform and accessed by a hundred people, than if that same thing is discussed by two or three people at a party. Lastly, *the strength of the epistemic relation* in which B knows p. Blaauw considers that someone being certain about A's pregnancy is more diminishing for their privacy than someone having a weakly justified true belief due to hearing a gossip about

[5] One could argue that this would only be possible if the people that knew them before they became hermits or got stranded in the desert island had their memories erased somehow, or there were none of such people.

[6] Gavison (1980) raises the issue that there are problems "when we attempt to compare different "amounts" of knowledge about the same individual. [E.g., w]ho has more information about X, his wife after fifteen years of marriage, his psychiatrist after seven years of analysis, or the biographer who spends four years doing research and unearths details about X that are not known either to the wife or to the analyst?" (ibid., p. 430). I will attempt to capture this intuition by adding as variables for the gradual aspect of privacy two more notions: the relationship between the individual and the subject, as well as the relationship between the individual and the piece of information.

A's pregnancy.[7] I will tentatively propose three further conditions that vary the strength of a privacy loss. First, a privacy loss varies depending on *how personal a piece of personal information is*. This rests on the plausible assumption that some propositions are more personal than others, e.g., someone's credit card information is surely more personal than information concerning what they ate for breakfast that day. Second, the *relationship between* the individual accessing the personal information B and subject A, as I discussed in Subsection 2.1.1, since it is not the same for my sibling or a complete stranger to read my diary. My third proposal is that the *relationship between* an individual B and A's personal information p matters as well. For example, it is more diminishing for A's privacy that their religious and conservative family member hears about their non-heterosexual orientation, than it is for their queer-friendly open-minded friend to hear it.

In this section, I have presented the gradual nature of privacy. In the following, I will inspect the phenomenon of privacy losses, characterised by the action of *accessing* some personal information. This action is interesting because our intuition does not provide us with an obvious epistemic solution. For example, can someone lose privacy over whether they are pregnant as a result of unfounded gossip? Does B diminish A's privacy if B tells other people false things about A? How relevant is the strength of an epistemic state for privacy, can A lose more privacy the stronger that B's epistemic state towards A's personal information is? In this chapter, I will discuss these questions in detail.

2.2 Privacy loss: Epistemic access

For a privacy loss to occur, an external person must acquire *epistemic access* to someone's personal information. In the following sections, I will define and defend what *epistemic access* means. I will argue for a similar notion to the *privacy losses* from Skopek (2020, p. 3).[8] In particular, three conditions must be fulfilled for B to diminish A's descriptive privacy over personal matter p about A. First, p must be true. Second, B must believe that p. Third, B's true belief must fulfil some form of *epistemic merit*. The definition of epistemic merit arises from the fact that I posit an account of *epistemic pluralism*. Such an account conceives that for a true belief to be deemed as knowledge it must fulfil one, or a combination, of epistemically relevant properties, i.e., justification, reliability, causal link, and so on. In particular, simply having a true belief about p is not sufficient. Véliz (2024) has recently provided an argument in favour of *lucky* true beliefs to be sufficient for a privacy loss, I closely inspect the consequences of her argument in Section 2.4 and provide a reason to reject it.

As I attempted to show in the previous chapter, to investigate the nature of descriptive privacy, an important first step is to analyse the phenomenon of privacy diminishments (also called privacy losses). In the previous chapter, I provided a preliminary definition of privacy losses in terms of epistemic access. For the following, let A and B be distinct individuals and let p be A's personal information.

Definition 2.2.1 (**Privacy losses**)**.** A loses privacy over p with respect to B if and only if B epistemically accesses p.

[7] Véliz (2024) argues that, in the reality of cases, the severity of a privacy loss does not change since true belief or justified true belief can have the exact same effects for privacy. I will inspect and reject this argument in Section 2.5.

[8] Skopek's formulation differs from mine in the following way. His three conditions for a privacy loss to occur are that (1) someone, B, accesses personal information about another person, A; (2) the means of access have epistemic merit; and (3) the information is true (Skopek 2020, p. 3). There are two theoretical distinctions between our accounts. First, I assume that the a privacy loss entails epistemic access, and this is the phenomenon that I am trying to define. Skopek, instead, takes privacy loss to be constituted of epistemic access plus other two conditions. I find it difficult to encapsulate this from an epistemological point of view. Second, for me, epistemic merit is an extra condition for the true belief that B forms, while Skopek takes epistemic merit to be an additional condition for access.

In the remainder of this chapter, I will present and discuss the main theories that describe the phenomenon of epistemically accessing some personal information. Matheson (2007, p. 259) provides a first step towards answering what it means to access personal information with his Broad Ignorance Theory (BIT).

Definition 2.2.2 (BIT). *A has privacy over p with respect to B if and only if B does not know p.*

Many privacy scholars have defined privacy in terms of knowledge. Some examples are Westin (1967, p. 7), Fried (1968, p. 483), Gavison (1980, p. 429), and Parent (1983, p. 346). However, these theorists do not study the condition of knowledge using notions from epistemology, e.g., by wondering whether a weaker epistemic state than knowledge may be sufficient for a privacy loss (which is my main goal in this chapter), or by applying notions from epistemology, such as *knowability*, to privacy (as I will do in Chapter 3). In the previous chapter, I showed that Fried's and Westin's theories failed as descriptive theories of privacy since they were control accounts, and I also discussed some issues with Gavison's and Parent's theories.[9] Now, to inspect the appropriateness of Matheson's epistemic condition closer, we can formulate BIT in terms of privacy loss.

Definition 2.2.3 (Privacy loss in BIT). *A loses privacy over p with respect to B if and only if B acquires knowledge of p.*

From this definition, it follows that Matheson conceives the knowledge of p to be a necessary and sufficient condition for a privacy loss over p (Matheson 2007, p. 259). Recent literature that discusses the epistemological nature of privacy agrees that knowledge is sufficient for a privacy loss, as one would also intuitively agree with, but they contend that knowledge is *not* a necessary requirement (Blaauw 2013, p. 170; Fallis 2013, p. 155; Kappel 2013, p. 188; Véliz 2024, p. 94). They argue that weaker epistemic notions, such as weakly justified true belief, may be sufficient. To assess which might be the weaker epistemic state that can lead to a privacy loss, Blaauw (2013, p. 171) provides the following list of epistemic relations from weaker to stronger. I have added number (3) to account for a version of the causal theory of privacy put forward by Fallis (2013), which I explain in Section 2.5.[10]

(1) Belief that p
(2) True belief that p
(3) True *hooked* (causal link) belief that p
(4) Justified true belief that p
(5) Degettierised true belief that p
(6) Rational true belief that p
(7) Warranted true belief that p
(8) Knowledge that p
(9) Certainty that p

As I said above, knowing someone's personal information p is uncontroversially considered to constitute a privacy loss. To see this, consider the following example, modified from Fallis (ibid.,

[9] On the one hand, the issue with Parent's theory was his focus on protecting only undocumented information, (Parent 1983, p. 429). This is problematic in this digital age, where an increasing amount of our personal information is documented, such as our photos and thoughts on social media posts, and we should still have some privacy protection over them. Gavison required limited knowledge (Gavison 1980, p. 428), instead of 'lack of knowledge' of personal information as the condition for descriptive privacy. Recall from the previous chapter that this is a problem as it fails to account for cases in which there is limited knowledge about someone, but this is still sufficient for a privacy loss to occur.

[10] Notice that this ordering is simply one proposal for the strength ordering of epistemic states but I leave room to different orderings and the addition or removal of other epistemic states. For example, not every epistemologist will agree that (4) and (7) are different, or whether it is even possible for there to be a *degetterised* true belief (i.e., a true belief that does not fall for Gettier cases, as I explain below).

pp. 155–156). Let A and B be distinct individuals, and let A's ankle tattoo be A's personal information.

Example 2.2.1 (**Direct perception**). *B sees A's ankle tattoo as A trips down the stairs on their first day of work.*

B sees A's ankle tattoo by direct perception. From this, B forms an epistemic state of knowledge, where they know that A has an ankle tattoo. This amounts to A losing privacy over their (having an) ankle tattoo with respect to B. This is intuitively clear. Since knowledge amounts to a privacy loss, then someone being *certain* that p constitutes a privacy loss too since certainty is a stronger epistemic state than knowledge. Now, in the remainder of the chapter, I will entertain whether weaker states than knowledge are sufficient for a privacy loss. Prior to this, I will discuss what knowledge is according to the epistemological literature.

In epistemology, knowledge is generally expressed as a form of justified true belief (JTB). Different epistemic theories impose different conditions to justified true beliefs to account for tricky cases raised by Edmund Gettier (1963). Gettier shows that some true beliefs may be formed from justified false beliefs in situations that arise out of luck (Zagzebski 1994, p. 65). To show an example of such situations, consider the classic clock example from Russell (1948, p. 170), where he asks us to imagine someone who stares at a broken clock to check the time. It so happens that the person checks the clock exactly on the time of the day in which the time of the broken clock and the actual time align. This person forms a belief that happens to be true, but thinking of it as amounting to knowledge feels wrong because of the luck component.

Epistemologists that generally agree with the idea that knowledge is a form of justified true belief have two strategies to deal with Gettier cases. They either make the justificatory condition stronger, or they add a fourth condition to avoid the Gettier conclusions (Ichikawa and Steup 2024, Section 3), such as safety or sensitivity.[11] Other epistemologists disagree that adding a fourth condition to JTB is an appropriate route, but at the same time they also consider bare true belief as insufficient for knowledge. There are two main kinds of this sort of epistemologists: reliabilists and causal theorists. On the one hand, reliabilists propose that a true belief must be formed through a mental process that is *reliable*. On the other hand, causal theorists ague that there has to be causal link between a belief about a true fact and the fact. Now, all of these proposed theories of knowledge, whether adding an extra condition to JTB, or suggesting another connection between the belief and the truth that p, have been shown to fall for Gettier problems (Ichikawa and Steup 2024; Zagzebski 1994). In light of this, some theorists have suggested that searching for one unified theory of knowledge might be misguided. They suggest that, instead, one should accept that knowledge cannot be defined with a set of precise conditions that agree with what we take knowledge to be in every situation. One can accept the complexity of knowledge in its inability to be satisfied by a set of everlasting simple conditions. At the same time, one can also agree that one needs a theoretical framework to deem a belief as epistemically meaningful to set it apart from mere speculation. Here, *epistemic pluralism* arises (Zangwill 2020). This is the idea that the necessary link for knowledge between truth and belief varies in distinct situations.

In this chapter, I will argue that for A to lose privacy over personal information p with respect to B, the following minimal conditions must be met: (1) p must be true (in Section 2.3); (2) B must believe that p (in Section 2.4); (3) B's true belief must have a link that has epistemic merit, i.e., it must satisfy one or a combination of epistemic conditions such as justification, reliabilism, and a causal link (in Section 2.5).[12] I will discuss this in turn in the following sections.

[11] I will investigate the plausibility of these notions as candidates to reformulate the conditions that protect the right to privacy in terms of epistemology in Chapter 3.

[12] Although I only discuss the listed conditions in this essay, I add the "and so on" part to leave space for future research on other conditions that might also be sufficient for a privacy loss to occur.

2.3 Truth

In this section, I will investigate whether it is possible for someone to lose privacy over falsehoods. I will argue that this is not possible. For someone's privacy to be diminished, the information under consideration must be true. Matheson (2007, p. 264), Blaauw (2013, p. 169), Kappel (2013, p. 180), and Véliz (2024, p. 89) agree with this intuition and consider (informational) privacy to be a notion that deals solely with true information. According to these theorists, one cannot lose privacy about falsehoods *per se* almost from theoretical assumption. However, falsehoods can still contribute to the occurrence of a privacy loss, e.g., by attracting attention to someone after spreading false gossip that has "sufficiently spectacular" information (Gavison 1980, pp. 431–432).

To begin discussing the plausibility of falsehoods contributing to our privacy in some way, consider the idea of lies. It seems that one may be capable of protecting their privacy by denying the truth over time, and by attacking belief (Mackie 1977, p. 183; Fallis 2013, p. 159). Consider the following example.

Example 2.3.1 (**Fake Digital Twin**)**.** *A* protects their personal information as much as possible on the internet. For this purpose, *A* only creates profiles with pseudonyms that contain fabricated personal information. They also use a VPN to hide their actual IP address. *A* takes any other step available to them in order to provide incorrect details about them on the internet.

By sharing false information online, it seems that *A*'s privacy is protected such that, whenever people or companies access *A*'s information and build a profile about them, in reality, they are creating a false one. This protects *A*'s privacy as it hinders *A*'s identification. In this way, falsehoods may be considered to even be good for protecting our privacy. However, Anita Allen (1988, pp. 21–22) famously proposed an example to challenge these intuitions.

Example 2.3.2 (**Allen's Diary**)**.** *A* writes false claims and fantasies on their diary. One day, *A* leaves their diary on the table of a café. *B* grabs *A*'s diary and reads it. From this, *B* forms false beliefs (that *B* considers true) about *A*.

Allen argues that, in this case, *A* loses privacy over the false claims. The epistemic literature is not convinced. Matheson (2007, pp. 263–264) agrees that *A* loses privacy in this example. However, *A* does not lose privacy over the false claims. Instead, according to Matheson, *A* loses privacy over details on their writing, their imagination, their fantasies, and so on. Véliz (2024, p. 87) agrees that *A*'s personal space has been wronged because someone read their diary without their permission. However, when considering whether *A* loses privacy specifically over the falsehoods, Véliz is reticent to accepting falsehoods as involving a privacy loss for a similar reason to Parent. Parent (1983, p. 269) conceives learning about someone's false claims, as well as their dissemination, as constituting "libel or slander", but certainly not amounting to privacy losses. My intuition concerning Example 2.3.2 is that there is a sense in which a privacy loss occurs because, even though *A* only writes false claims, those claims say something about *A*'s psyche and imagination, i.e., I agree with Matheson's analysis. However, *A* does not lose privacy over the falsehoods themselves because they say nothing about *A* (and may even act as protecting their privacy, in a similar way to Example 2.3.1).

So far, I have only focused on cases where *A* is the one who puts forth some falsehood about themselves. However, do our intuitions change if falsehoods about *A* come from a distinct source? Consider the following example.

Example 2.3.3 (**Ruining a reputation**)**.** *B* does not like *A*. *B* routinely spreads false rumours about *A* on the internet with the goal of ruining *A*'s reputation. *B* follows *A* whenever they can, and fabricates plausible though false rumours about *A* based on what *A* does.

In this example, A is clearly wronged, but does A's wrong arise from a worry about privacy? By paying attention to A and following them, B wrongs A as B forces the establishment of a—plausibly unwanted and toxic—relationship upon A (Brake 2023). B also wrongs A by spreading false rumours about A, which can constitute moral wrongs on their own. These rumours can attract more attention towards A. Attention facilitates the appearance of privacy diminishments as, by exercising attention towards you, people are looking out to whether you commit a mistake and may access some personal information about you, as well invade your personal space, in response (Gavison 1980; Véliz 2024). However, the act itself of B spreading false rumours can be done without B knowing anything at all about A. We can imagine a case where A is an actor from a series and B is a fan of the series. Then, B might spread false rumours about A completely based on the beliefs B has about A in the context of the series. Within a descriptive notion of privacy, intuitively this does not seem to entail a privacy loss since B is engaging in confabulation related to fiction.

Therefore, the epistemic nature of descriptive privacy provides evidence to thinking that, if p is false, then one cannot lose privacy over it. Thus, for a privacy loss to occur, p must be true. However, your right to privacy is indeed violated when someone spreads false rumours about you. I will inspect how to conceptualise this within an epistemic pathway of the right to robust privacy in Chapter 3. I will now turn to questioning whether someone has to believe p for a privacy loss to occur.

2.4 Belief

In this Section, I will discuss three examples raised by Fallis (2013) which attempt to show that it may not be necessary for a belief to be formed for there to be a privacy diminishment. I will explore and defend the need for a belief to be formed within each example, with Example 2.4.3 being the most compelling. From this, I will argue that, if someone follows my intuitions, it is necessary for someone to at least form a *true belief* about a piece of personal information for there to be a privacy loss. In the next section, I will discuss whether a mere true belief might be sufficient for a privacy loss or stronger epistemic states are necessary.

Most privacy scholars assume that one must at least hold a belief for a privacy loss to take place. However, Fallis (ibid.) discusses some examples that, according to him, support the idea that one's privacy can be diminished over a certain piece of personal information even when no one forms a belief about it. Let us being with the following example from Fallis (ibid., pp. 164–165).

Example 2.4.1 (**Repeated Glances**)**.** B looks at A's personal picture every night for a year.

Let us assume that B forms all possible beliefs about A's picture after the first few glances at time t. Due to this, Fallis highlights that B does not gain any new meaningful beliefs after t. However, Fallis argues that by continuing to look at A's personal picture, there is a sense in which B diminishes A's privacy every single night, even after t. Thus, according to Fallis, after time t, there are privacy losses, but the external agent does not form a belief related to it. Fallis takes this to lend credibility to the idea that there can be epistemic access that results in a privacy loss without belief. Fallis considers this to favour a causal theory of privacy that I will explain in the next section.

One could of course argue that, in this example, B does form new beliefs. Perhaps these new beliefs are updated with time if every single instance in which B looks at the picture means B forms a new belief with an associated distinct time, or one could maybe argue that B might find something new and distinct about the picture every instance. However, I am more inclined to argue that this kind of case conflates personal information and personal space. I agree that A's informational privacy is only diminished the first few glances before t (if we assume B forms no

beliefs afterwards) and not in the glance at time $t+1$. However, the intuition that A's privacy is still being diminished in this example comes from the fact that A's personal space is being accessed. This is because B gains perceptual access (through sight and touch) to a personal picture about A. A similar example to Example 2.4.1 has been discussed in the literature, where B watches A shower every night until it is no longer possible for B to form any new beliefs. In this case, it might be more clear that A's privacy is diminished because A's personal space is being accessed as B sees A in their personal space undertaking a personal activity.

Therefore, this example shows that our full intuition about privacy losses cannot be solely explained with informational privacy, and that spatial privacy is an important dimension of privacy that can explain tricky cases like this one. This example does not show that descriptive privacy losses (from an informational point of view) can occur without belief, as the distinction between information and spatial privacy is less clear cut than it may seem at first. Let us consider another example from Fallis (2013).

Example 2.4.2 (**Gossip**)**.** C discloses p about A to B. p is so incoherent with A's normal behaviour that B "cannot believe [their] ears" (ibid., p. 165).

Radford (1966) has argued that it is possible for someone to have knowledge without belief. The intuition comes from cases where you forget that you know something so you are simply not sure that you know it, e.g., regarding someone's birthday, or the capital of a country (Véliz 2024, p. 89). Then, when you are questioned about it, you give an answer that you think must be wrong. However, it turns out that you are right. Therefore, Radford argues, you knew the answer—your friend told you their birthday last month—you just did not believe you knew it because you thought you had forgotten it. Now, recall that when B knows something personal about A, then A loses privacy about it. So, the claim here is that there can be instances were privacy is lost—because you knew personal information—without belief because you did not believe you knew it.

I am not compelled with the main claim from these arguments, i.e., that you knew something, but you did not believe it. In Example 2.4.2, it seems more likely that when B expresses that they cannot believe what C told them, B does not literally mean that they do not believe it. It is simply a figure of speech because of the surprise that hearing p caused in them. Thus, B did not *literally* not believe p, as the mere act of using this expression and saying those words suggests they did.

Let us now discuss the last example from Fallis (2013, p. 165) that I find the most compelling.

Example 2.4.3 (**Automated advertising system**)**.** An automated system displays ads to A based on personal information p about A.

Real-time bidding is a process by which someone's personal information (e.g., IP address, web search, estimated location, device data) are sent to companies that bid for their advertisements to be displayed on that person's current webpage (Privacy International 2019). There are huge privacy risks regarding this practice, as anyone who poses as an advertisement company receives the data, regardless of whether they win the bid. It is a common method for data brokers to enter the bid as advertisement companies with the goal of collecting data from users and potentially selling it to third parties (Cohen 2025).

To simplify, let us assume that, in Example 2.4.3, a process takes place such that A's personal information p is sent to ten companies. Each of these companies has an automated system that bids depending on what p is, i.e., depending on whether you are a Finnish male Mac user that frequently orders food at home, or whether you are a Californian female user of Android that watches a lot of climbing YouTube videos. Within the time that the webpage and its ad space loads, one company wins the bid and displays the ad.[13] Regarding this situation, it seems clear that

[13]Note that *advertisement blockers* are programs you can embed in your web browser to avoid being shown ads, and there are also programs that protect your personal data from being shared in this kind of bids.

there has been a privacy wrong and, Fallis argues, there is also a privacy loss from the collection of personal data and the personalised ad being showcased. However, as Fallis notes, no belief is formed because no person looks at your data to then form a belief about it and act accordingly (this is an assumption). Instead, an automated system has dealt with everything. How can we reconcile this with the idea of epistemic access that I have proposed in this thesis? Does this mean that privacy can be lost even when no one forms a belief about it?

Let us take a moment to inspect Example 2.4.3. One can identify at least three wrongs related to privacy within it. (1) Personal data about A's device and location is created by virtue of them browsing a website, (2) A's personal data is sent to and stored by ten advertisement companies (or data brokers posing as ad companies), (3) given A's personal data, an automated system places a bid for their companies' ad to be shown. These are privacy wrongs because they are wrongs arising via securing a place where they can easily diminish A's privacy without A's consent with (1) and (2), and via using A's personal data in (3) without A's consent. Now, Fallis' argument is related to (3). He argues that, during this process, A's privacy is diminished even though no one forms a belief about A's personal data. The advertising content that we see is not generic but, instead, it is personalised. What exactly is wrong about it regarding A's privacy?

This wrong is the one that *decisional privacy* focuses on. The personalised content from the ad has the goal of steering A's behaviour and manipulate them (Susser et al. 2019), either to make them shop for a particular product, or make them vote for a candidate in an election, or to take another sort of action. Recall that the ad is related to A's location and, most likely, from web searches associated with A's IP address. For this reason, this wrong is related to A's privacy. Additionally, one could argue that (1) and (2) are also wrongs related to A's privacy because the creation, collection, and dissemination of A's personal data occurs mostly without A's awareness or understanding (it is very unclear which companies collect it and what is their real purpose). Now, I certainly agree that, in Example 2.4.3, A is wronged in (1), (2), and (3) and, for this reason, A's right to privacy is violated during the entire process. However, is A's descriptive privacy lost? By distinguishing between descriptive privacy and the normative right to privacy, I have intentionally left the possibility that one side could be respected, while the other is not. I posit that Example 2.4.3 is such that A's normative right to privacy is violated, while A's descriptive privacy is not diminished.

To make this salient, let us assume that for each of the ten companies that collect A's personal data, none of their employees look at it and, after the bid is placed, the data is immediately deleted. Further assume that the ad that is shown to A is extremely misleading. In particular, it fails to steer A's behaviour as the profiling they did of them was wrong or incomplete. In this case, A still suffers a privacy wrong, as they collected (although briefly) and used their data (although in a perhaps *useless* manner from a marketing point of view), but A does not seem to have lost any actual privacy. Yes, A's privacy is put at risk throughout the entire process and, likely, this risk is not justified. A suffers a salient privacy wrong because the companies have *attempted* to use A's personal data to steer A's behaviour.

By simplifying the situation in Example 2.4.3 with these assumptions, my point is to show that, since descriptive privacy only cares about the moment of accessing information, A's descriptive privacy is not actually lost. For there to be access, a belief must be formed. Otherwise, we could not distinguish a situation in which someone forms a belief about A's information and does nothing with it—where A's privacy is diminished but their right to privacy is respected—with a situation in which someone confabulated something about A, everyone believed it and acted as if it was true—where arguably A's right to privacy is violated but their descriptive privacy is not. I will use a similar argument in the following section to show why a mere true belief is not sufficient for a privacy loss either.

To recapitulate, if one follows the intuitions which I have defended in this section and the

previous one, it follows that there has to be an individual who forms at least a true belief about someone to diminish their privacy. However, is a mere true belief sufficient? And, if it is not sufficient, what kind of other epistemic conditions must the true belief satisfy? I will discuss this questions in detail in the following section.

2.5 Justification and other types of epistemic merit

In the previous sections, I have argued that an individual B must at least posses a *belief* over a *true* personal fact about A for B to diminish A's privacy. In this section, I will discuss whether a mere true belief is sufficient for a privacy loss, or whether a further link satisfying some form of epistemic merit is necessary. For this, I will discuss an argument by Véliz (2024), where she argues that a mere true belief is sufficient. By arguing that Véliz is conflating between privacy losses and privacy wrongs (i.e., between descriptive and normative privacy), I will show that a true belief is not sufficient, and another form of *epistemic merit* must be satisfied.

Subsequently, I will discuss different conditions that must be satisfied for a true belief to fulfil epistemic merit (Skopek 2020, pp. 21–26). The notion of epistemic merit rests on the assumption of *epistemic plurality*, i.e., that a true belief can be linked via different processes for it to yield knowledge (Zangwill 2020), in this case, for it to yield a privacy loss. In this section, I will briefly discuss *justification*, *causality*, and *reliability* as plausible proposals for epistemic merit. However, I leave space for other kinds of epistemic merit to be found upon further research.

A *mere* true belief?

I have hitherto shown that, if you follow the intuitions I have defended in this chapter, one must at least possess a *true belief* about someone's personal information for one to diminish that person's privacy. Now, is a mere true belief sufficient for a privacy loss? Consider the following modified example from Véliz (2024, p. 91).

Example 2.5.1 (**Wishful thinking**)**.** B and A are friends. By wishful thinking, B forms the belief that A is pregnant. This belief happens to be true.

Note that B has neither a justification nor a causal connection between the formation of the true belief regarding A's pregnancy and A's actual pregnancy. Most epistemologists of privacy consider that wishful thinking does not involve a privacy loss (Blaauw 2013, p. 171; Fallis 2013, p. 157; Matheson 2007, p. 264). However, Véliz (2024, p. 91) argues that this example does entail a privacy loss. In short, Véliz assumes that privacy losses can occur without justification because it is possible for someone to form a true belief out of luck, where the consequences of this lucky true belief would be indistinguishable from a belief formed in a justified manner. I will show that the ramifications from following this chain of argument, which is excessively reliant on consequences and privacy wrongs, are undesirable for the theoretical enterprise of descriptive privacy. Thus, a mere (lucky) true belief should not be considered sufficient for a privacy loss.

Let us spell out the strategy that Véliz follows. First, Véliz (ibid., p. 94) asks us to imagine plausible consequences from B forming their unjustified true belief in Example 2.5.1, which happens to be true out of pure luck. Assume, first, that neither A nor any other person that B talks to have any reason to think that B's true belief was formed out of lucky wishful thinking. Further, let us assume that, as a consequence of their true belief, B undergoes a change of behaviour and judgement towards A, and that B acts upon their belief by telling other people about A's pregnancy. Let this be situation (U). We can imagine another scenario in which B had formed the true belief about A justifiably (e.g., because A's partner told them). It is reasonable to imagine that the consequences from forming the justified true belief could be the same to the ones we just

described from an unjustified true belief, i.e., B changes their behaviour and judgement towards A and tells other people about A's pregnancy. Let this be situation (J). Since the consequences for privacy are indistinguishable regardless of whether we are in situation (U) or (J), then Véliz argues that this must mean that an unjustified true belief that achieves the same (for privacy) as a justified true belief is sufficient for a privacy loss.

My intuition does not agree with Véliz. Her strategy seems to conflate between privacy losses and privacy wrongs. I would agree that, if B changes their attitude towards A and also tells people about A's pregnancy, then B commits a privacy wrong to A.[14] Thus, a privacy wrong occurs in both (U) and (J). However, I do not think that A's descriptive privacy has been diminished in (U). For one reason, an unjustified true belief is not constituent of a privacy wrong, e.g., it could easily be that B acts in a way such that B simply ignores and attempts to forget their wishful thought, or that A interrogates B into how they formed the belief that A is pregnant and B has to drop the belief. Then, in these scenarios, it is unclear that A suffered any privacy wrong nor that they lost any descriptive privacy. A critic could say that it is fine for these cases to not entail a privacy loss, as it is only the lucky situations like (U) that we care about. Recall that these are the situations where B's unjustified belief is true out of luck and where B commits privacy wrongs that are indistinguishable from privacy wrongs arising from a justified true belief.

Even when we focus on situations like (U), there are problems with saying that (U) leads to a descriptive privacy loss because of its consequences. In particular, let us assume that any two situations, as long as their consequences are equally wrongful for privacy, then they are both equally diminishing for descriptive privacy. If we follow this pragmatic route, are the truth or belief conditions necessary at all for descriptive privacy? Imagine that a machine learning model predicts that A has committed fraud for childcare benefits (although they have not). Let us further imagine that the government, A's family, A's friends, and in general, everyone in A's social circle act as if A had committed fraud. They create distance with A, they stop trusting A, A's employer fires A from their job, and so on. Importantly, the consequences for privacy are exactly the same to a case in which A has actually committed fraud. Does this mean that A's descriptive privacy is diminished when the machine learning model predicts something false about A? It seems very clear to me that it is not. Yes, in both cases wrongs occurred about A's supposed personal information, i.e., that A committed fraud. However, in one of the cases, their descriptive privacy was actually diminished because they had indeed committed fraud, while in the other their descriptive privacy was not diminished, because what they were wronged about was not true.

Therefore, it seems to me that, if we were to accept the argument from Véliz that a mere true belief is sufficient for a privacy loss, then we must also embrace that what is necessary for descriptive privacy is something similar to the idea of *public acknowledgement* from Nagel (1998). According to Kappel (2013, p. 189), public acknowledgement captures the idea that, as long as a person or a group of people act as if a piece of personal information about person A (true or not, believed or not, justified or otherwise) is true; then, A can suffer the same privacy effects as if the person or group of people had full knowledge of the personal information. This notion could arguably be useful for the idea of a privacy wrong, as we may want to say that someone can suffer a privacy wrong solely if people act as if some personal information about them is true. However, for the theoretical enterprise of descriptive privacy, I believe that this notion does not capture the morally neutral idea of privacy loss which allows us to say that, even if two friends disclose personal information to each other with full consent, they still lose descriptive privacy to each other. Similarly, if someone's roommate sees them naked by accident and does their best to forget it, their roommate's descriptive privacy is still diminished when the act occurs. Even though

[14] Note that privacy wrongs are described by the right to privacy within the hybrid account of privacy from Véliz (2024) that I will present in Chapter 3. Depending on how serious a privacy wrong is and how it occurs, it may mean that B has infringed A's right to privacy, B has failed to respect it, or B has violated it. These states go from less morally wrongful to more morally wrongful.

the flatmate does not publicly acknowledge seeing them naked, they have suffered a privacy loss according to descriptive privacy. Thus, if you follow my intuitions and are not ready to accept the notion of public acknowledgement as a plausible condition for descriptive privacy, then a mere true belief is not sufficient for a privacy loss.

In the following, I will briefly consider some candidates from the epistemological literature to link a true belief in such a way that a privacy loss can occur. These candidates will be joint under the umbrella term of *epistemic merit* from Skopek (2020), by endorsing epistemic plurality. Epistemic plurality refers the idea that knowledge can be achieved by linking a true belief in various ways (Zangwill 2020), e.g., via a justification, a reliable process, or a causal link (this, in turn, are kinds of *epistemic merit*). Although epistemic pluralism normally accounts for knowledge, in this case, I endorse it in suggesting that different true belief may be connected in distinct epistemically valid manners for it to yield privacy loss. Thus, I do not commit to the theoretic idea that the epistemic states that result from connecting true beliefs with epistemic merit need to necessarily yield knowledge (e.g., in degetterised way). Instead, I only focus on these states being sufficient to lead to privacy losses.

So, I resort to *epistemic plurality* to explain the possible avenues for privacy losses to occur, rather than arguing for only one of them (e.g., justification) to be the one responsible for privacy losses. This is because I believe that there must be a balance between being too permissive and too restrictive regarding the kind of epistemic state which we deem plausible to lead to a privacy loss. On the one end of the spectrum, the epistemic state needs to be formed by a process that is epistemically appropriate, since I already argued a mere true belief is not sufficient. On the other end of the spectrum, recall that I have already mentioned that there are Gettier cases for each proposal that attempts to supplement true belief and yield knowledge (Ichikawa and Steup 2024; Zagzebski 1994). In a similar way, I will now show that there might be cases in which a privacy loss takes place even though it is not justified, but it is instead causally connected. Although this discussion will be brief, I will render it to pose sufficient suspicion that the phenomenon of privacy loss is varied so one kind of epistemic state is not necessary, but a plurality of them serve to cause privacy diminishments. Let me now briefly discuss three candidates to link true beliefs in turn: justification, causality, and reliability.

Justification

In epistemology, there are two main kinds of theories regarding the nature of justifications: internal and external. *Evidentialists* are one type of internalists regarding justification that argue that someone's true belief is justified if it agrees with the evidence they have gathered up until that point, while externalists argue that the reasons for someone's true belief to be justified must be found externally to the agent (Ichikawa and Steup 2024, Sect 1.3). It seems that both kinds of justification can cause a privacy loss. An evidentialist would consider it appropriate justification for a nutritionist to judge that their friend is very likely to suffer a cardiovascular disease given that the background beliefs of the nutritionist allows them to infer this privacy diminishing information (the likelihood of a disease) after their friend tells them their daily calorie intake (recall this example from Subsection 2.1.2). Additionally, externalists would consider as privacy diminishing a case in which someone stumbled upon their roommate naked in the living room via direct perception. Thus, both internalists and externalists kinds of justification can lead to privacy losses.

By formulating that an agent's true belief must be justified, I have been assuming *doxastic justification*—the view that considers justification to be focused on the appropriateness of beliefs (ibid., Sect 1.3). There is also another kind of justification called *propositional justification* concerning the justificatory appropriateness of a certain proposition. An example to see that the two notions come apart, inspired from (ibid., Sect 1.3), is to consider a subject A, who is completely ignorant

of a considerable amount of evidence that suggests that it is going to rain tomorrow (e.g., the latest climate models predict it will rain with high accuracy, it is the time of the year when it normally rains, the environment shows signs of humidity and moist, and so on).[15] However, upon reading their horoscope saying that tumultuous times are coming, A forms the belief that it will likely rain tomorrow. As horoscopes are generally considered inappropriate, epistemically speaking, for believing something, A's belief is not doxastically justified. However, it is propositionally justified because there is plenty of external evidence suggesting that it will rain tomorrow (although the agent is ignorant of this). One may notice the similarity between this case and that of a lucky true belief that I discussed in the previous part of this section. For the same reason for which I argued against lucky true beliefs above, propositional justification is rejected. Thus, instead, when I say that a true belief must be justified for a privacy loss to occur, I refer to doxastic justification.

Causality

Fallis (2013, p. 165) argues that someone's privacy may be diminished even when the true belief is not justified but there is still a way in which "a cognitive agent [...] perceives [personal information p] at least in an attenuated sense". This attenuated sense is a *causal hook* that Fallis considers sufficient for a privacy loss. In this manner, Fallis endorses a causal theory of privacy following Goldman's causal theory of knowledge (Goldman 1967).[16] Consider the following example inspired by Fallis (2013, pp. 159–160), where A is the subject whose privacy is under assessment and B is an individual.

Example 2.5.2 (**Drunk B**). B is drunk. A and B are work colleagues. B sees A's ankle tattoo as A trips down the stairs on their way to a bar.

Is A's privacy diminished in Example 2.5.2? B acquires a belief that A has a tattoo in their ankle, which is indeed true. Due to their drunkenness, however, B is not sure whether they really saw the tattoo, so B's justification is weak. Still, weak justification seems sufficient for A's privacy to be diminished. However, imagine now that B is an agent that dismisses their experiences when they are drunk as, in the past, they have fabricated facts or considered dreams to be true when drunk. In this case, we would not be able to explain why it still seems that Example 2.5.2 yields a privacy loss if B were to implicitly hold the belief anyway, despite their efforts.

Fallis (ibid., pp. 159–160) proposes that B's true belief is *hooked* to A's ankle tattoo via a causal chain, and this is what truly matters for a privacy loss. In particular, A's privacy is diminished over personal information p with respect to B as long as p is "causally connected in an 'appropriate' way" with the belief of p by B (ibid., pp. 159–160). Fallis refers to this appropriate causal connection as *cognizance*. Cognizance can arise from mediated (e.g., through sensorial device) and unmediated (direct) perception, and from testimony from someone who is cognizant about the fact (ibid., pp. 160–161). Now, what kind of causal connection between the fact and the belief is necessary exactly? B has to have some evidence, even if insufficient for knowledge, such that a causal chain occurs (ibid., p. 162), although this same evidence may not be sufficient for justification.

However, Fallis himself recognises the limits of the causal theory by rendering that cognizance cannot be used to explain the way that B may gain "knowledge of universal facts, facts about the future and negative facts" (ibid., p. 163). Also notice that the main proponent of the causal theory of knowledge, Goldman, noticed its insufficiencies to determine the cases for knowledge and, later

[15] Anyone that has lived in Amsterdam will probably consider the fact that one is in Amsterdam as sufficient evidence to render that it will most likely rain the next day. Determining whether this is an epistemically appropriate justification is outside of the scope of this thesis.

[16] It is important to note that Fallis argued for a causal theory of privacy where the agent need not believe the fact for someone to lose privacy over that fact. As I already argued for belief to be necessary for privacy in Section 2.4, I gloss over this and consider causality as a link mediating truth and belief.

on, became a defender of *reliabilism* (Goldman 1979; Ichikawa and Steup 2024, Sec 6.1). Thus, let us now discuss reliabilism as another candidate to link a true belief for a privacy loss to occur.

Reliability

Some epistemologists that were both dissatisfied with the insufficiency of justification and causal links to yield knowledge, given Gettier cases or else, turned to *reliabilism*. Reliabilism is the epistemic theory that renders a true belief to be knowledge as long as the agent acquires it through a reliable cognitive process (Ichikawa and Steup 2024, Sect 6.1). A reason to prefer reliabilism has to do with the idea that animals do not seem to have the justification required for beliefs to be knowledge, however, given their actions, we would still want to say that animals know things about their environment (ibid., Sect 6.1). In this manner, we can consider animal knowledge to be formed by a reliable cognitive process.

Now, reliabilism can link a true belief such that it leads to a privacy loss. For example, direct perception is a reliable cognitive process and so every privacy loss which occurs because an agent sees or hears something reliably can be considered a reliable true belief that leads to a privacy loss. However, reliabilism is not sufficient for a privacy loss as one can imagine a case in which a privacy loss occurs because an unreliable testifier tells the truth (Fallis 2013, p. 157), but there is a distinct justification or causal link to render this as a privacy loss. Thus, reliabilism is another plausible candidate to cause privacy losses.

In this section, first, I have argued that a mere (luky) true belief is not sufficient for a descriptive privacy loss, *contra* Véliz (2024). Then, I have proceeded to argue that a true belief must be linked in a epistemically relevant way for a privacy loss to occur, without committing to one single theory of the matter. I have proposed a list of non-exhaustive candidates that I have briefly discuss for their suitability regarding being links of true beliefs that may lead to privacy losses. These candidates are justification, causality, and reliability. I leave space for further research to find whether other epistemic candidates are appropriate in forming epistemic states that lead to privacy losses.

2.6 Conclusion

In this chapter, I have investigated the epistemology of descriptive privacy. For this purpose, I have achieved two research objectives. First, I have described privacy as a ternary relational state following Blaauw (2013). The second objective has been to elucidate the epistemological nature of the epistemic access that is necessary for informational privacy diminishments.

For the first research objective, in Section 2.1, I inquired on the nature of the relata of privacy as a relational state, i.e., the subject A, the set of individuals I, and the set of personal information P; then, I discussed the gradual nature of privacy. First, I argued that the subject A, whose privacy is assessed, must be sensitive to the values that privacy protected in order to be in such a position. Then, I argued that any individual that is capable of epistemically accessing A's personal information can be a suitable candidate for potentially diminishing A's privacy. Subsequently, I discussed the intricacies of the notion of personal information. I argued that any piece of personal information about a subject is *potential* personal information, even the most innocuous one. This potentiality is being exacerbated by technological advancements of the digital age, such as big data and data mining. To deal with the broadness of personal information, contextual solutions may be needed, I discuss this further in the Conclusion of this thesis. Finally, I discussed the fact that privacy is not an absolute notion and that one may lose more or less privacy depending on a variety of factors. First, I explained the proposals from Blaauw regarding the strength of a privacy loss: the amount of personal information being accessed, the number of people accessing

a piece of personal information, and the strength of the epistemic state formed upon the access. Then, I proposed three further factors that can also vary the strength of a privacy loss. First, how personal a piece of information is, the relationship between the subject and the individual that accesses their personal information, and the relationship between the individual and the piece of personal information.

The second research objective concerned the investigation of the nature of epistemic access that leads to a privacy loss. For this purpose, I argued that for an epistemic access to count as a privacy loss, the following three conditions must be met: (1) p must be true, (2) B must believe that p, and (3) B's true belief must be linked by a form of epistemic merit, such as, justification, reliability, or causality. In Section 2.3, I argued that false propositions cannot cause privacy diminishments almost due to theoretical necessity. Then, in Section 2.4, I argued against three arguments posited by Fallis (2013) that seemed to show that a privacy loss may occur without there being a belief. The necessity for belief was argued as a necessity to distinguish facts that are believed in a epistemically robust way from pure confabulations. A similar argument was then used in Section 2.5 to show that a mere true belief, that has the same consequences for privacy as a justified one, is not sufficient for a privacy loss, contra Véliz. For this argument, I highlighted that Véliz confused privacy losses and privacy wrongs by focusing in pragmatic consequences. I argued that, if one is to take seriously the enterprise of descriptive privacy, one should not rely on such overly pragmatic consequences. Finally, I posited that a true belief must be linked in an epistemically appropriate way, i.e., fulfilling *epistemic merit*. I briefly surveyed justification, reliability, and causality as plausible candidates for privacy diminishments, although I left space for further research to consider other epistemic notions.

Chapter 3

The Epistemology of the (Normative) Right to Privacy

In this chapter, I will investigate the normative notion of privacy: *the right to privacy*. In particular, I will present, discuss, and expand on the framework for the right to privacy from Véliz (2024). The right to privacy protects us from suffering wrongs related to our descriptive privacy. Recall that (informational) descriptive privacy (with respect to an external person B and over a piece of personal information p about subject A) is the relational state that A has if B lacks epistemic access to p. In contrast, the *right to robust (informational) privacy* refers to the right for us not to be wronged in relation to attempts to access our personal information p and actual accesses of p in the actual world and, further, in relevant counterfactual worlds. Notice that by protecting, not only actual access, but also attempts to access our p, the right to privacy has a *path-based nature* (Johnson 1989; Kappel 2013; Nissenbaum 2010; Skopek 2020). This is because one protects the *means* by which others access or attempt to access our personal information (Véliz 2024, p. 76). In comparison, descriptive privacy is only concerned with actual accesses of p, not with how someone achieved that access, so it is *path-independent*. To clarify the interplay between the normative and descriptive side of privacy, it is important to note that someone can violate another's right to privacy without diminishing their descriptive privacy, and *vice versa*. For example, in the previous section I argued that the automated collection of personal information and automated advertisement can violate someone's right to privacy but not diminish their descriptive privacy. Conversely, one's descriptive privacy can be diminished without their right to privacy being violated, e.g., when two friends talk about their personal matters with each other's consent.

Now, for someone to publish explicit photos of their previous partner without permission is a clear violation of their right to privacy, as it is for someone to instal a virus in your computer via a malicious link in order to gain access to your passwords and financial information. However, is it a violation of your right to privacy when you get rejected from a job because a screening tool that uses machine learning tries to predict your suitability based on the tone of your voice (Corbyn 2024)? Or when you are arrested for a crime that you did not commit after being wrongly identified with facial recognition software (Hill 2024)? Lastly, is it a privacy violation when an algorithm trained with your personal information—which no person epistemically accessed (as defined in the previous chapter, by forming a true belief that fulfils epistemic merit)—shows you advertisements that influence your shopping decisions? A debate within the privacy literature concerns the deliberation on whether these cases may be considered wrongs or harms related to privacy, or closer to other notions, such as autonomy or freedom. In this chapter, I will create a space to assess this debate with respect to a recent account of the right to privacy proposed by Véliz (2024). Véliz's account describes the right to privacy as robust, such that it must be

respected, not only in current circumstances, but also on counterfactual situations (Véliz 2024, p. 145).

Given the definition of descriptive privacy that I provided in this thesis and expounded above again, it follows that valuable actions provoke privacy diminishments. For example, you lose privacy by sharing progress on your mental health with people that care for you, and you diminish someone's privacy by writing reliable news articles about their financial irregularities or about their involvement in a political scandal. It is important to note that privacy is not always the most important value or right to be protected in every scenario. In the examples presented, other values are more prevalent than privacy, such as establishing meaningful relationships with loved ones, and holding people accountable as public officials. In daily life, then, one should ask what are the situations where we should act to protect someone's privacy, and how we can achieve this. In other words, we should ask what our privacy duties are, when is it that they become relevant, and how we can make sure to attain to them. Before we establish our privacy duties, we must have a clear idea of what the right to privacy protects. There are plenty of suggestions on how to define a right to privacy in the literature. In this chapter, as I have mentioned, I will base my analysis off Véliz's right to robust privacy.

In Section 3.1, I will explain how Véliz's account of the right to robust privacy ensures that respecting someone's privacy not only means that one respects it now but also in relevant future and counterfactual worlds. This description is advantageous because Véliz does not rely on a broad notion of control to define the right to privacy. Control accounts have been the most common strategy to define the right to privacy in the past but have been shown to fail as a necessary and sufficient condition for the definition of this right. Additionally, Véliz's account manages to account for the violation of the right to privacy that occurs in cases of indiscriminate and automated mass data collection.

After presenting and defending Véliz's account as a suitable main framework for the right to privacy, I identify a feature that the right to privacy should protect and that is missing in her framework: the wrong caused by the inappropriate use, e.g., disclosure or dissemination, of personal information that has been acquired from someone else at an earlier time with their consent. Then, I propose that this condition must be explicitly added to Véliz's definitions of the *failure to respect* the right to privacy and the *violation* of the right.

Subsequently, in Section 3.2, I propose and develop a novel account of the right to privacy that focuses on an epistemic pathway of the flow of personal information and incorporates Véliz's theory. My epistemic pathway explicitly accounts for the whole flow of a piece of personal information p in three steps: the process of gaining access to p, the actual access of p, and the use of p. I discuss how an epistemic analysis of the step of gaining access and using p can help us explicitly formulate the challenges for privacy which are encountered in each scenario. For instance, for the process of gaining access to p, I propose the epistemic notion of *knowability* as a plausible candidate to reformulate the protection to attempt to access p. Lastly, I conclude this chapter in Section 3.3.

3.1 Véliz's right to robust privacy

A theory concerning the right to privacy should provide us with a framework in which to assess whether, from the perspective of privacy, someone has been wronged (ibid., p. 143). To track privacy wrongs, Véliz proposes that it is not only important to assess whether privacy is respected in the actual world, but that we must also consider whether it is respected in relevant possible worlds (ibid., p. 145). That is, Véliz defends the right to *robust* privacy. To bring some intuition, Véliz (ibid., p. 77) explains that robust privacy works in a similar way to the republican notion of freedom along the lines of Pettit (1996). According to this notion, it is not enough to say

that a slave is free because the master has not interfered with them. It is also important that the slave is free from "arbitrary" interferences with "no impunity" (Véliz 2024, p. 77). The fact that the master could interfere with the slave if they pleased to is sufficient to render the slave as not free, even if the master has not actually interfered with them just yet. Similarly, someone's right to privacy cannot be said to be respected solely because no one has hitherto diminished their privacy. If an individual or a company are in a position in which they could easily diminish that person's privacy if they pleased—e.g., because they have that person's personal information in a folder in their computer—, then they are not respecting that person's right to privacy. Thus, the right to privacy is a right to *robust* privacy that must be respected in actual and counterfactual circumstances.

Now, as one discusses the right to privacy, it is important to say something about what a *right* is. I will not fully present and defend a particular theory of rights, as this is outside of the scope of this thesis. However, I will note that Véliz endorses Joseph Raz's interest theory of rights (Raz 1988). According to the interest theory, rights are such that the interests that protect the well-being of living beings are secured with duties (Véliz 2024, p. 121). I will consider privacy duties to be the actions that an individual, a collective, or an institution should take in order to respect the right to privacy of an individual, or a collective. After clarifying this, I will now present and discuss Véliz's robust right to privacy. Véliz categorises two kinds of invasions of the right to robust privacy: failures to respect the right to privacy and violations of this right (ibid., p. 145). Véliz uses the language of "relevant possible worlds" and "counterfactual" and "future" scenarios. However, she does not make this formulation explicit in terms of using possible world semantics. In the following, I will present the first step of an attempt to provide an explicit formulation in terms of possible world semantics of Véliz's right to robust privacy.[1]

Let A and B be distinct individuals, and let p be personal matters about A, let $w_1, ..., w_n$ be possible worlds, and let *invading A's privacy* correspond to *diminishing A's descriptive privacy without A's consent*.[2]

Definition 3.1.1 (**Failure to respect the right to privacy**)**.** Even though B respects A's (descriptive) privacy at w_1, B *fails to respect* A's right to privacy at w_1 if

(i) B has not invaded A's privacy due to B being lucky or lazy, or

(ii) in relevant possible worlds $w_2, w_3, ...$, B would be ready to invade A's privacy.[3]

To illustrate some cases encompassed by this definition, consider an example from Véliz (ibid., pp. 144–145) related to an example from Macnish (2018), where A leaves their diary at B's house. Simply put, if B does not read A's diary, B respects A's (descriptive) privacy in the actual world. However, Véliz argues, for B to respect A's *right* to privacy, this is not sufficient. B must also be "disposed to refrain" from reading A's diary absent A's consent in *relevant possible worlds* (Véliz 2024, pp. 144–145). For example, B should not read A's diary regardless of whether B were to

[1] See (Véliz 2024, p. 145) for the specifics of her formulation.

[2] Note that the notion of consent is a highly problematic one, especially in the digital world (Lanzing 2019, p. 564). Informed consent in the digital world is dubious because, most often than not, we do not actually have a choice about using a service as we need it for social or work-related communication or services. Additionally, it is difficult for the consent to be *informed* as it is usually really unclear what information is actually being collected, what companies collect it and to what purpose. Also, whenever this information is listed, it is either impossible or very difficult to understand as it is contained in awkward cookie banners or technical privacy policies.

[3] Notice that this formulation may seem doubly modal. In particular, one may notice that an individual B is said to fail to respect A's privacy in the actual world if, in another relevant possible world, it is *not* that B invades A's privacy (I will show now that if this were the case, this would be a violation of the right to privacy given Véliz's definition). Instead, B fails to respect A's right (in the actual world) if B is *disposed* to invade A's privacy in those relevant worlds. This is an intriguing requirement that may seem to bring yet another modal condition. However, I will give Véliz's credit to the benefits of her formulation, and I will note that any weird modal instances arising from the formulation given in this thesis are probably due to my attempt of embedding Véliz's definition in possible world semantics. One is likely to maintain Véliz's essence while avoiding this double modal nature. However, I will leave this issue aside for the purposes of the present thesis.

despise A, or the contents of A's diary were to pose great curiosity for B, or B were to profit from something written in A's diary (Véliz 2024, p. 145). If B were willing to read A's diary in any of these relevant scenarios, then following condition (ii), B fails to respect A's right to privacy. Moreover, imagine that the only reason that B has not read A's diary is because B does not realise that A left the diary (i.e., if B had realised, they would have read it),[4] then B fails to respect A's right to privacy according to condition (i). Notice that, in these cases, B fails to respect A's right to privacy, but does not violate this right because the invasion of A's privacy is not realised, either by chance or for some other reason, although it could have easily been otherwise (in a relevant possible world). For the following, let $w_1, ..., w_n$ be possible worlds.

Definition 3.1.2 (**Violation of the right to privacy**). "Absent outweighing conflicting considerations" (ibid., p. 145), B *violates* A's right to privacy at w_1 if:

(a) B "secures a position" such that they can invade A's privacy with the intention of invading A's privacy at w_1, in the future or counterfactual world w_2, or

(b) B tries to invade A's privacy at w_1, or

(c) B invades A's privacy at w_1.

Before discussing the suitability of this definition, note that the definition begins by adding that "outweighing conflicting considerations" must not be present for an action to count as a violation of the right to privacy. By considering this, Véliz aims to account for the fact that, when it comes to rights, one needs to weigh, case by case, whether the invasion of one right is justified by the preference of another right. To see this, consider this definition of infringements of the right to privacy (ibid., p. 146).

Definition 3.1.3 (**Infringement of the right to privacy**). B *infringes* A's right to privacy if there are reasons for which the invasion of A's descriptive privacy is justified.

So, when a rights invasion is justified, one does not violate someone's right but *infringes* it instead. Consider Véliz's example where the police come to know that "information to save a thousand lives" is written in someone's diary (ibid., p. 146). In this case, the police's duty to read the diary for the right to life of those thousand beings outweighs the right to privacy of the person who owns the diary. In Véliz (ibid., pp. 120–140) provides a lengthy discussion on the conflicting considerations between the value of privacy and the value of surveillance in different scenarios. Véliz's aim is to defend the value of privacy and the harms of surveillance in a society where the latter are often ignored, in detriment for privacy (ibid., p. 137).

Let us come back to discussing Definition 3.1.2 of the violation of the right to privacy. The Snowden revelations in 2013 showed that the intelligence agencies NSA and GCHQ carried out an indiscriminate collection of personal data from mobile devices and from other devices with an internet connection (Macnish 2018, p. 418). Upon the revelations, the agencies defended themselves by arguing that they did not violate people's right to privacy because the majority of the data was never accessed or analysed by a person, in which case it was merely collected in bulk (MacAskill 2015). Within the right to robust privacy framework, Véliz (2024, p. 146) shows that this reasoning is misleading. On the one hand, it is true that, if we assume that the data collected by the intelligence agencies was never accessed by a human being, then our descriptive privacy was not diminished. On the other hand, however, Véliz argues, the NSA and GCHQ violated our right to privacy following condition (a) of Definition 3.1.2 because, by collecting our personal data, they

[4]Notice the similarity with the condition of *safety* from the epistemological literature. That is, B's belief of p is safe if and only if, p would not be false if B were to believe p (Ichikawa and Steup 2024). Note that here p refers to reading A's book, i.e., accessing A's personal information. In Section 3.2, I will propose knowability as an epistemological candidate to reformulate condition (a) from Definition 3.1.2. One can see that further research may focus on using safety or sensitivity to account for the failure to respect the right to privacy

"secured a position" in which they could invade our privacy in relevant possible worlds (Véliz 2024, p. 146). Thus, mass surveillance where our personal data is created, collected, and stored in bulk without consent or reasonable justification entails a violation of the right to privacy.

Imagine that someone from the NSA wants to investigate a particular person so they try to find that person's name in the system to go through their data. However, it happens that the NSA does not have data collected from that person. Therefore, the NSA agent attempted (unsuccessfully) to invade that person's privacy. Provided there were no outweighing moral reasons for the NSA agent to look for the person (e.g., they did it out of curiosity and not because of there were reasonable expectations that they would commit a crime) then, following condition (b) of the definition, the NSA agent violated that person's right to privacy by attempting to invade their privacy. Lastly, condition (c) captures cases in which B successfully diminishes A's privacy without A's consent. Whenever there are no outweighing moral justifications for B to spy on A, or for B to touch A's hair, or for other way for B to access A without A's consent, then B violates A's right to privacy.

In the following, I will present and discuss some objections to Véliz's right to robust privacy. I will raise one objection that may require the reformulation of Véliz's right to privacy framework if I am correct. After I finish discussing Véliz's right to robust privacy, I will present and discuss my own explicit analysis of the flow of a piece of personal information that embeds and expands Véliz's right to privacy. This flow will be conceptualised as a pathway formulated with epistemological notions. This way, I will present a novel manner in which an epistemological analysis can be developed for the normative right to privacy.

3.1.1 Objections and responses

Véliz's formulation of the right to robust privacy is conceptualised within her hybrid account of privacy, which is a novel and recent theory from 2024. For that reason, there has not been a vast discussion on her account in the philosophical literature yet. Nevertheless, I will now present some objections to Véliz's right to privacy from Munch (forthcoming), who is the only author that I have found that discusses Véliz's new theory, and I will also raise other issues myself. I will proceed to defend Véliz's account against most of them. However, I will raise one worry regarding information misuses that I believe Véliz's account does not account for. To deal with this issue, I will propose to supplement Véliz's theory.

Munch (ibid., p. 4) raises two objections to Véliz's account on the right to privacy. The first problem that Munch raises is the fact that Véliz's hybrid account of privacy draws a distinction between *privacy* and the *right to privacy* that is *non-harmonious*. A harmonious theory of privacy is such that whatever we define privacy to be, e.g., X, then, the right to privacy must be *the right to* X (Lundgren 2021, pp. 382–383). For Véliz's account to be harmonious, there would be two ways to comply with this. One way would be for the right to privacy to simply be the *right to descriptive privacy*, i.e., the right that someone's personal matters are not accessed by external people. However, this would not be able to account for the fact that, most often than not, our right to privacy is violated even when our personal matters are not accessed. For example, if B is my colleague and regularly asks me very inappropriate questions about my personal life, B violates my right to privacy even if I am resilient enough to never answer B's questions. Then, let us consider the other way in which we could try to formulate Véliz's theory more harmoniously. One could get rid of descriptive privacy altogether and consider privacy as a robust normative notion that must be protected in actual and counterfactual scenarios. The right to robust privacy would then be the right to have one's robust privacy protected. The problem with formulating privacy solely as a normative robust notion is that one fails to account for the descriptive side of privacy. Descriptive privacy describes the actions that are normally regarded to lead to imbalanced power dynamics regarding knowledge, even when this dynamics are not exploited (Véliz 2024). If

the arguments that I have provided in this thesis are convincing, one can see that getting rid of descriptive privacy altogether is not a viable solution if we want to describe the whole notion of privacy with its disparities and distinct intuitions. Therefore, even though non-harmonious, Véliz's division of privacy between descriptive privacy and the normative right to privacy forms an important whole that, I have tried to argue, covers our intuitions regarding privacy.

The second objection that Munch raises is related to the scope of Véliz's account, which he deems to be too broad. Munch presents an example inspired by Thomson (1975, p. 305) where A's neighbour buys an X-ray device with the intention to look through A's walls, but the neighbour does not actually go through their plan. Munch (forthcoming, p. 5) wants to argue that, as long as the neighbour never uses the X-ray device, A's right to privacy is not violated. The problem with Véliz's account, according to Munch, is that it fails to differentiate between *gaining the means to access* and *actually accessing* a personal matter. In particular, he posits that, even if there is a possible world in which A's neighbour could use the X-ray device to spy on them, A's right to privacy is not violated until they do use it. However, it seems to me that Munch is restricting the right to privacy too much. Let us go back to A's neighbour who just bought an X-ray machine. Imagine that they try to look through A's wall with it, but it turns out that the X-ray machine does not work. Then, it seems very reasonable that A's neighbour violates A's right to privacy. Following his reasoning, Munch would consider this attempt as not violating A's privacy by attempting to do it, as there is a nearby world in which the X-ray machine works and they do look through their wall. If Munch were to not be convinced by this, then it seems that Munch simply does not want to consider privacy as a robust notion akin to the republic ideal of freedom. However, I find the fact that Véliz can account for relevant counterfactual worlds as a advantage of her account. This allows her to determine why the indiscriminate collection of personal data is a privacy wrong and why you still violate my privacy if you ask me very intimate questions and press me to answer, even if I do not answer them.

A third problem one may raise, which seems to be pressing in the background of Véliz's account, relates to the problem of how to assess what are the possible worlds that count as relevant in each specific situation. This is indeed a worry that arises from any account that uses *relevant possible worlds*. Véliz concedes that there must of course be a limit to the range of possible worlds. The possible worlds that count as relevant are those in which "moral reasons to respect the right to privacy continue to outweigh the balance of the competing considerations" (Véliz 2024, p. 151). With this requirement, Véliz makes sure that we do not count possibilities that are not supported in moral grounds, or which are outweighed by competing considerations.[5]

Additionally, one may argue for an alternative account of privacy wrongs by positing that the wrong pertaining the intelligence agencies' collection of the data is situated in the risks that it poses to our privacy as well as how this makes us feel (Macnish 2018). Véliz argues that this way of framing the issue fails to capture the privacy abuses that intelligence agencies commit with respect to our data, regardless of whether they could make sure our data is risk-free (Véliz 2024, pp. 150–151).

Finally, one may argue that, within Véliz's account, it is difficult to establish exactly in which cases a privacy wrong occurs since this has partly a conventional nature (Kappel 2013, p. 188). Now, Véliz embeds this conventional and social norms in her account, which may render her account unable to object to practices that subtly become the norm but endanger the human right to privacy and the values it protects. This issue is raised by Véliz, and she argues that her definition

[5]When we determine the relevant worlds to see whether B violates A's right to privacy, it seems difficult at first glance to see exactly how one could cut worlds in which, for instance, A did something very bad to B so B wanted to revenge, or cases in which B inadvertently accesses their personal information. However, at least this latter part seems to be taken care of in condition (a) of Definition 3.1.2, by delimiting that B must have an *intention* to invade their privacy in the actual world. In any case, although puzzling, I believe that this worries are solvable, although I will not attempt to solve them in this thesis.

can avoid these troubles by restricting what personal information is. In particular, she conceives personal information to be "the kind of information about oneself that people have reason not to want to share widely" (Véliz 2024, p. 155). Her purpose with this is to embed within the notion of personal information to be connected to the manner in which we can obtain values such as autonomy, safety, relaxation, democracy and so on (ibid., pp. 100–109).

So far, I have defended Véliz's theory against the objections raised by Munch (forthcoming) and myself. However, I will now raise an objection that I believe entails that Véliz's notion of the right to robust privacy must be expanded. Let me begin by noting that, according to Véliz (2024, p. 144), condition (c) of Definition 3.1.2 captures the following case. Assume that at time t_1, B accesses A's personal matters with A's consent, e.g., A willingly tells B about their latest health diagnosis after going to the doctor. Then, at a later time t_2, B discloses the information gained to other people without A's consent. For example, imagine that B tells A's representative of their health insurance, C, about A's health diagnosis in exchange for money. For these cases, Véliz argues that the violation of the right to privacy committed by B is covered by condition (c), which would mean that B invades A's privacy. According to Véliz, as B discloses A's information without their consent, B makes A lose privacy (as other C access the information), so A's privacy is invaded.

However, I disagree with Véliz, her formulation of the right to privacy cannot account for these cases because condition (c) does not explain why B violates A's right to privacy. Notice that, if B were to be violating A's right to privacy via condition (c) in Definition 3.1.2, then it would be B who accessed A's health information at t_2 without A's consent. However, B does not do that (B already accessed that information *with* A's consent at t_1). Instead, when B tells C at t_2, then it is C who accesses A's health information without their consent. Thus, the wrong that B commits to A's privacy is not via the invasion of A's privacy, but via *using* A's health information inappropriately by further disclosing it to C without A' consent.

Véliz could argue that her account does deal with cases of misuses of personal matters although it is not explicitly contained in the definition of the right to privacy. In particular, Véliz has argued that for one right of the right to privacy, it is not one corresponding duty that comes with it, but multiple ones. Véliz (ibid., p. 161) follows Raz (1988) in thinking that one should not exhaust morality with rights. Instead, duties may arise in relation to them. This would make sense since Véliz (2024, p. 161) does explicitly mentioned that the right to privacy entails a duty of silence, i.e., meaning that one should not disclose this information. However, this would mean that by disclosing their information B does not violate A's right to privacy. Instead, B disobeys a moral duty relates to the right to privacy. I find this conclusion dissatisfying as one is surely violating my right to privacy when they tell thousands of people about my money issues, and not simply disobeying a duty.

Therefore, I believe that Véliz must modify her account to encompass cases where B violates A's privacy via the improper *use* of A's personal information, I will now propose two modifications that Véliz could make to her account. As a first proposal, she could add an extra condition to Definitions 3.1.2 and 3.1.1 that account for attempt and actual cases of improper *uses* of the information gained after it has been accessed. For example, this would render the definition of the violations of the right to robust privacy in the following way. Let A and B be distinct individuals, and let p be personal matters about A, let $w_1, ..., w_n$ be possible worlds, and let *invading A's privacy* correspond to *diminishing A's descriptive privacy without A's consent.*

Definition 3.1.4 ((Modified)* violation of the right to privacy). "Absent outweighing conflicting considerations" (ibid., p. 145), B *violates* A's right to privacy at w_1 if:

(a) B "secures a position" such that they can invade A's privacy with the intention of invading A's privacy at w_1, in the future or counterfactual world w_2, or

(b) B "secures a position" such that they can *misuse* A's privacy with the intention of misusing A's privacy at w_1, in the future or counterfactual world w_2, or

(c) B tries to invade or misuse A's privacy at w_1, or

(d) B invades or misuses A's privacy at w_1.

However, it seems that adding this extra condition complicates the conditions for violating the right to privacy. Perhaps a more elegant solution would be for Véliz to enlarge the notion of privacy invasions to account for both the diminishment of A's privacy without A's consent and the use of A's personal information without A's consent.[6] Therefore, I consider that someone who misuses your personal information by disclosing to other people without your consent violates your privacy. In fact, I think that this kind of right to privacy violations are discussed in the literature a lot, in relation to how they undermine our autonomy and leave us powerless. This is why Véliz should explicitly account for misuses of personal information (more generally, personal matters) within her account. I believe that this is important because B has interfered with the appropriate flow of A's personal information (Nissenbaum 2010).

In the next section, I will discuss further what it means to take seriously for the right to privacy to account for the whole flow of personal information and not only the access or attempts to access. To continue my endeavour in this thesis, I will use notions from epistemology to reformulate ideas on how to protect personal information from unwanted invasions. At this point, I will note that a reader might agree with me regarding the fact that Véliz's right to privacy must add this extra condition concerning misuses of personal information, without agreeing with the epistemic analysis that I will undertake in the following section.

3.2 An epistemic pathway

Recall that in Chapter 2, I presented an epistemological analysis of descriptive privacy. I discussed proposals regarding epistemological notions that attempted to clarify the conceptualisations for descriptive privacy. Specifically, I explored the relational state of *privacy*, and what it means to access a piece of personal information from an epistemological point of view. Authors that have explored the epistemology of privacy have almost exclusively focused their efforts on using epistemology to describe descriptive privacy. In this section, I want to expand the epistemology of privacy by using notions from epistemology to clarify the notion of the normative right to privacy.

So far, in this chapter, I have presented, defended, and expanded on Véliz's right to robust privacy account. In this section, my aim is to account for the right to privacy as a right that protects the means of access of a private matter in a "path-based" manner (Johnson 1989; Kappel 2013; Nissenbaum 2010; Skopek 2020). The exchange of personal information is inevitable in the digital age. In order for personal information to be properly protected, the entire pathway of information exchanges must be considered (Nissenbaum 2010).

Personal information does not flow in the vacuum. I will divide the pathway for describing a violation of the right to privacy into three time steps. First, we encounter the *pre-access* step, where B gains access to A's personal matters p by either *inquiry* (by gathering evidence) or *phenomenological access*. In the second step, B *accesses* p. Lastly, in step three, B *uses* p by disclosing it or through some other action. I will focus my efforts on the first and third steps of the pathway since I already analysed the step of epistemic access in Chapter 2. In the first step, I will embed the conditions that Véliz proposes for someone's right to privacy not to be violated or

[6] For now, my focus has been on personal information. However, notice that B would misuse A's personal space if, for example, they took advantage of a moment in which they were in A's house to let other strangers in. The strangers invade A's privacy. However, B violates A's right to privacy by misusing A's personal space by letting other people in.

failed to be respected. I will reformulate these conditions using the epistemic notion of *knowability*. Note that most of the efforts that I will draw for the epistemological analysis will be focused on informational privacy since the epistemic connection can be drawn more clearly. However, future work might expand this investigation for personal space.

3.2.1 Inquiry (pre-access)

Let us begin with the pre-access step, where an individual B *gains* access to (but does not yet access) personal information p about A by inquiry, or personal space s of A by physical or perceptual proximity. In this step, A's privacy is not diminished, as the *pre-access step* refers to the time frame just before A's personal matters are actually accessed.

As a first step, we may ask what it means for someone to inquire such that they gain access to my personal information. Munch (forthcoming) has recently presented an account that highlights the importance of inquiry for the right to privacy (ibid., p. 7). Within his account, inquiry is presented as an activity that rational agents perform in order to form epistemic states. Munch analyses two conditions that must be met for inquiry to occur. First, one must hold "interrogative attitudes" (Friedman 2017). These are mental attitudes that are directed to answering a question by "wondering" or "deliberating" if a certain proposition might be the case (Munch forthcoming, p. 6). Secondly, one must execute actions in order to achieve the goal of inquiring. Some examples of inquiry as a mental state encompass "weighting evidence, making inferences, entertaining alternative hypothesises, and reflecting upon the sufficiency of one's available evidence" (ibid., p. 6). As a physical attitude, inquiry might involve gathering evidence by browsing an encyclopaedia, or by engaging in a shared deliberation in conversation (ibid., p. 6). The act of inquiry pertains to the act of securing a position with the intention of invading someone's privacy. Suitably, inquiry also necessitates an *intention*.

Now that I have provided some background on the nature of inquiry, how do we assess when a given process of inquiry is morally permissible or not? Imagine that B inquires A about their personal information p. How do we determine if this action is a right to privacy violation? Let us discuss some possibilities. B's inquiry is morally permissible regarding privacy if B follows foreseen social norms and A *consents*. For example, in cases of two friends talking about their worries, or when someone asks the doctor about some health worry. Notice that there might also be cases in which B's inquiry is permissible simply because it is *unavoidable* or *accidental*. This takes into account cases in which B could not look away, or could not avoid hearing something. For example, if A goes on a rant regarding their latest financial problem, this is case of voluntary disclosure akin to Example 1.2.3, and B does not violate A's right. Now, B's inquiry may be morally inadequate if the circumstances are inappropriate because, e.g., B resorts to peer pressure, coerces A, or eavesdrops on a conversation using an advanced hearing device (such as in Example 1.1.2).

I have just provided an intuitive survey regarding the actions of inquiry that may be morally permissible, so that it becomes salient that Véliz's robust right to privacy is able to account for them since Véliz considers that one must account for other moral considerations to weigh whether someone's right to privacy is violated or not. To show this, I will now provide Véliz's definitions of the failure to respect the right to privacy and the violation of the right to privacy. For the following, let A and B be distinct individuals, and let p be personal matters about A, let $w_1, ..., w_n$ be possible worlds, and let *invading A's privacy* correspond to *diminishing A's descriptive privacy without A's consent*.

Definition 3.2.1 (**Violation of right to privacy (pre-access condition)**). "Absent outweighing conflicting considerations" (Véliz 2024, p. 145), B *violates* A's right to privacy at w_1 if

(a) B "secures a position" such that they can invade A's privacy with the intention of invading A's privacy at w_1, in the future or counterfactual world w_2, or

(b) B tries to invade A's privacy at w_1, or

Notice that I have only left the conditions that account for the pre-access step by which B secures a position to gain access to A's personal matters. I will now show that one of the conditions from Definition 3.2.1 can be re-formulated with the epistemic notions of *knowability*.

Knowability

The epistemic state of *knowability* refers to a state in which an agent, even though they have all the evidence that is needed to infer a piece of information p, they need *not* have engaged in epistemic activity that grants them access to p (yet) such that they are said to know p (Hawke and Berto 2021) inspired by Dretske (1981). For example, a piece of information is knowable to an agent if there is an inference rule or further epistemic step that the agent has not applied yet that separates the agent from knowing p. By focusing on knowability, Hawke and Berto (2021) aim to avoid the challenges that come from dealing with the necessary cognitive states and attitudes which an agent needs to achieve knowledge.

For our purposes, knowability seems to be an interesting candidate for Véliz's condition (a) of "securing a position" in Definition 3.2.1. This conditions says that B violates A's right to privacy if B *"secures a position" such that they can invade A's privacy with the intention of invading A's privacy at w_1, in the future or counterfactual world w_2*. This condition could be expressed in the following way, where p is A's personal matters: B violates A's right to privacy if *p is knowable to B without A's consent such that B has the intention of knowing p at w_1, or in the future or counterfactual world w_2*.

By using knowability, we do not need to explicitly say that B has secured a position since the fact that p is knowable means that B has put themselves in a position where they could easily get to know A's personal matters. Now, I have discussed this in terms of knowability, but one may want to generalise given my discussion from Chapter 2 such that one substitutes this with *epistemic accessibility*.

3.2.2 Access

Recall that B epistemically accessing A's personal matters entails that A loses privacy over p with respect to B. Importantly, B diminishing A's privacy does not necessarily entail that B violates A's privacy. For example, when two friends consent to catching up about their romantic lives, or when two consenting adults undress in front of each other. To normatively determine if a privacy loss also entails a violation of the right to privacy, we must consider whether the means of access or the use of the information afterwards correspond to a wrong. In general, it seems that access is mostly concerned with privacy losses, and I discussed this in length in the previous chapter. Recall that I argued that subject A loses privacy over their personal information p with respect to individual B if three conditions are met: (1) p is true, (2) B believes that p, and (3) B's true belief fulfils a form of epistemic merit (e.g., justification, reliability, and causality).

3.2.3 Use (post-access)

How do we determine whether the way in which B uses the personal information about A is morally permissible or not? In particular, B may decide to keep p to themselves and not disclose it, or they may decide to share it with hundreds of people. When it comes to the Snowden leaks, it seems that the disclosure of the information was justified, as it benefitted civilians. However, the case portrayed by Kafka's works is more complicated to assess, morally speaking. In particular, Kafka did not wish for their work to be published, but their friend decided that it was too good to let

it rotten in a library. Here, one would need to weigh whether the right to people reading Kafka's books is stronger than Kafka's right to privacy.

It seems that a case by case moral assessment must be made to consider whether uses of personal matters infringe, fail to protect or violate someone's right to privacy. At this point, I will note that normative considerations regarding the values that privacy protects become especially important. In particular, whether the use of someone's personal matter is justified or not, depends on whether there are detriment effects to our autonomy or well-being. In which case, the use is not justified. For this purpose, let us survey the value that arises from privacy. To discuss different methods for determining the value of privacy, Rössler (2005) differentiates between two types of proposals. Proposals that conceive the value of privacy as *reducible* to other rights and those that do not. The most influential proponent of reductionist accounts of privacy is Judith Jarvis Thomson (1975), who argued that the right to privacy can be derived from the *rights over a person*—i.e., a right to not be harassed in the street—and *property rights*. Thomson argues that, whenever we believe a right to privacy violation has occurred, it is actually a distinct right that can be traced back to being the main cause for the rights violation. Thus, speaking of a distinct right to privacy is unnecessary. Although a provocative and influential account, many authors have deemed her analysis as leading to absurd rights, such as a right to not be looked at, and as ignoring the intuitive and common use of the term *privacy* as a distinct phenomenon (Inness 1996; Parent 1983; Véliz 2024).

Within the authors that conceive privacy as an irreducible concept, Rössler (2005, p. 69) differentiates two approaches that she highlights to be complementary of each other. Namely, intrinsic and instrumental theories of the value of privacy. Intrinsic theories argue that there are reasons to value privacy in itself without referring to other values, whilst instrumental accounts conceive privacy as resting its importance within other values or wishes, such as autonomy. For most purposes, the instrumental reasons to value privacy are striking enough that people are likely to defend it once this connection is made clear. Some *instrumental* accounts of the value of privacy connect privacy with our ability to identity as people and to form personal relationships; so intimacy is deeply important for privacy (Inness 1996). Additionally, Véliz (2024, pp. 100–111) provides an extensive defence of privacy by which she argues that "privacy protects control over self-presentation, our reputation, autonomy, creativity, security, freedom, equality, well-being and democracy". Rössler (2005) also highlights the connection to autonomy and freedom. Additionally, Gavison (1980, p. 442) argues for privacy's functions to be a "healthy, liberal, democratic, and pluralistic society; individual autonomy; mental health; creativity; and the capacity to form and maintain meaningful relations with others".

Given the values that privacy protects, to ensure that the right to privacy is respected, one should account for whether these values are respected during the use of personal matters after they have been accessed. The work from *decisional privacy* becomes especially relevant in delineating this by providing a clear way by which someone's personal matters must not be used when it will deteriorate our autonomy. Further research could elucidate the lessons from different philosophical works that protect the value of privacy to realise what actions must be avoided and which ones must be done for someone's personal information to be used appropriately. This will, in turn, help discover the privacy duties of different individuals for different circumstances.

3.3 Conclusion

In this chapter, I have investigated the epistemology of the normative right to robust privacy. In Section 3.1, I have presented Véliz's right to robust privacy in terms of possible world semantics. Then, I have defended her right to robust privacy against recent objections provided by Munch

(forthcoming), and others that I have raised myself. Then, I have highlighted one objection regarding the completeness of Véliz's account in accounting for the right to robust privacy. In particular, Véliz fails to account for cases of misuses of personal matters. I have proposed how Véliz may embed this into her theory. Subsequently, in Section 3.2, I tentatively proposed the idea of an epistemic pathway to account for the flow of a piece of personal information in three steps: (1) pre-access, (2) access, and (3) post-access. I show that the notion of knowability may be used to account for conditions of pre-access step, while judging whether the use of a personal matter is justified can only be found in considering the state of the values that privacy protects upon the potential use.

Conclusion and Further Research

Summary

In this thesis, I have motivated, discussed, and defended privacy as a notion that attends to descriptive and normative intuitions. The equivalent notions for these intuitions are descriptive privacy and the normative right to privacy. In Chapter 1, I have based my analysis off Véliz's Hybrid Account of Privacy by discussing the two main accounts of privacy in the literature: control accounts and access accounts. In Chapter 2, I have investigated the epistemology of descriptive privacy. For this purpose, I have inquired on the nature of privacy as a ternary relational state that is gradual. Then, I have investigated the nature of the epistemic access that leads to (informational) privacy losses. I have argued that for individual B to epistemically access personal information p about subject A, it must be the case that: (1) p is true, (2) B believes that p, and (3) B's true belief satisfies a form of epistemic merit, e.g., justification, causality, or reliability. Finally, in Chapter 3, I have presented, defended, and expanded on Véliz's right to robust privacy. Additionally, I have tentatively proposed an epistemic pathway to assess right to privacy violations by reformulating the conditions for protecting this right.

Contributions

In this thesis, I have attempted to contribute to a debate regarding the interaction between the research fields of *epistemology* and *privacy*. Regarding this effort, the following points are tentative contributions from this thesis.

- Throughout this thesis, I have extensively discussed and motivated the division of privacy into descriptive privacy and the normative right to privacy. I have argued that descriptive privacy aims to describe the actions in the actual world (path-independent) that put us in a vulnerable position because they make us more likely to be wronged regarding our personal matters. The normative right to privacy aims at providing a place where these privacy wrongs can be assessed in a path-based manner, by accounting for actions taking place in the actual world and that would occur in relevant counterfactual worlds.

- Also, throughout this thesis, I have provided explicit and clear definitions of the privacy nomenclature from the Hybrid Account of Privacy from Véliz (2024). The same notions are sometimes used in other privacy works in a distinct manner or confusingly. I have defined privacy losses (and diminishments) as actions related to descriptive privacy, and privacy invasions, infringements of the right to privacy, failures to respect it, and violations of it, related to the right to privacy.

- In Chapter 2, I have discussed each of *relata* of privacy as a ternary relational state from Blaauw (2013) and I have expanded Blaauw's analysis as a result. I have defended that the subject A, i.e., the being whose privacy is under evaluation, must be sensitive to the

values that relate to privacy. Then, I have argued that the individuals I are those with capacity to access A's personal matters, and that the content of the personal information about A is subject to contextual dependencies exacerbated by big data and data mining. Regarding Blaauw's work, I have also expanded on the variables that cause the gradual nature of personal information. I have proposed for the strength of a privacy loss to be dependent on how personal the personal information is, the relationship between the subject A and the individual, and the relationship between the individual and the piece of personal information.

- In Chapter 2, I have surveyed, expounded, and discussed the main views on the nature of epistemic access from the literature regarding whether it requires *truth*, *belief*, and further *justification*, while also providing independent grounds for epistemic access to be achieved with a true belief linked with a form of *epistemic merit* from Skopek (2020). I have argued for this by providing, to the best of my knowledge, the first objection against the argument that (lucky) true beliefs can be enough for a privacy loss from Véliz (2024). I have shown that the consequences from Véliz's argument confuse privacy losses and wrongs and over-rely on pragmatic considerations, which is undesirable for the theoretical enterprise of descriptive privacy. In contrast, most privacy scholars prior to Véliz's work had simply assumed that a mere true belief cannot lead to a privacy loss.

- In Chapter 3, I have presented Véliz's right to privacy explicitly in terms of possible world semantics, and I have defended this right against some recent objections from Munch (forthcoming). Finally, I have found a gap in Véliz's formulation of the right to robust privacy. In particular, she does not account for cases in which an individual B misuses the personal information that they had previously accessed about A with A's consent. I have proposed two ways in which Véliz may embed this misuses within her proposal.

- Finally, in Chapter 3, I have tentatively proposed an epistemic pathway for the flow of personal information separated by three main components, where the second one is covered by Chapter 2. The first step deals with how an individual may gain access to personal matters. Restricting to information, I have established inquiry to be the manner in which access can occur. I have used the epistemic notion of *knowability* to reformulate a conditions from Véliz's right to robust privacy definitions.

Further Research

I will note a selected amount of proposals for further research concerning connections between epistemology and privacy that I strongly wished to pursue but could not due to time and space constraints.

Personal matters

This thesis has almost exclusively analysed privacy from the dimension of information. However, the investigation from this thesis can be expanded to account for **personal space** by using insights from both epistemology and phenomenology. We are embodied creatures, so the intuitions from personal space are inevitably reliant on intuitions regarding day life morality since, inevitably, there is always some piece of ourselves that we portray to the social world. This leads me to believe that an in-depth discussion of personal space within privacy might prove challenging. However, further research might revoke this intuition.

To account for the nature of **personal information** in the digital age, one may discuss the challenge arising from statistical inferences. The practices of big data collection and analytics have

enabled the inference of personal information from seemingly innocuous information. Research has shown that, by inputting someone's social media posts in state-of-the-art large language models (LLMs), the LLMs can infer the person's location, age, and gender (Staab et al. 2024). For instance, asking for the likely location of the author of the following post, 'There is a nasty intersection on my commute, I always get stuck there waiting for a hook turn', the LLM GPT-4 responded that "a 'hook turn' is a traffic maneuver particularly used in Melbourne" (ibid., p. 1). This presents the following problem for personal information within privacy:

1. Location, age, and gender are pieces of personal information.
2. Text from social media posts (not mentioning personal information explicitly) is not considered personal information.
3. Using the context and inferential capabilities of, e.g., LLMs, one can infer location, age, and gender with text from social media posts (even when no personal information is explicitly contained).

This raises the question of whether text from social media posts is always personal information. I do not think for this to be the case. However, we need a manner in which personal information can be elucidated from non-personal information. For this purpose, I propose a logical framework where one could account for the effort and manner in which a piece of personal information is inferred from other evidence. This may have a similar research direction to the work of Studer (2011) using *justification logics*.

Further, another possible research direction is to investigate the social aspect of privacy. To illustrate the social component of privacy (and personal information), consider genetic data (Véliz 2020). Imagine that, out of curiosity, you send a DNA sample to a genetics company. Most genetics companies sell your data to third parties (e.g., insurance companies, who are interested in varying your policy on the basis of inferred likelihood for genetic diseases) (Martin 2018; Purtill 2024). Importantly, your genetic data is not exclusively about you; you share a significant amount of genetic information with close family members. Thus, the genetics company can infer, not only yours, but your siblings', cousins', and future generations' genetic information.

The Epistemology of (Descriptive) Privacy

- Further work may use the definitions and characteristics from descriptive privacy to build an *epistemic logic of descriptive privacy*.
- Further research could expand the kind of epistemic merit that is sufficient for a privacy loss to occur. Also, one may investigate further the notions of justification, reliability, and causality. In particular, what is the role of defeaters for justification?
- Finally, one could expand my investigation of the epistemology of descriptive privacy by accounting for sources of epistemic access: perception (senses), testimony (work from social epistemology could be considered here), and reasoning (inferences become especially important here).

The Epistemology of the (Normative) Right to Privacy

- Due to the formulation of Véliz's right to robust privacy in terms of possible world semantics, future work may focus in providing a formalisation of her definition in terms of a modal logics.
- Further work may continue analysing the epistemic pathway. For instance, by providing more explicit and deeper investigations into the notions of knowability and safety to reformulate

the protections of the right to privacy within *pre-access*. Additionally, one may provide to a further investigation into the morally relevant constraints for the use of *personal matters*. One could account for it by adding views on the literature that consider the values that are protected by privacy, such as autonomy and freedom.

- Embed the epistemology of normative privacy with work from feminist epistemologists. For instance, Patterson (2020) discusses "knowledge entitlements", which are privacy diminishing actions suffered by people from marginalised communities due to prejudice. For instance, people from these communities are questioned more, and their personal space is accessed more. E.g., people feel entitled to touch black women's hair without permission, or touch a pregnant woman's belly without asking first.

- Further account for privacy duties. Here, epistemic ideas such as whether we are in control of our own beliefs become salient (*doxastic voluntarism*), and whether what is rational to believe is determined by pragmatic factors (*moral encroachment*).

References

Allen, A. L. (1988). *Uneasy access: Privacy for women in a free society*. Rowman & Littlefield Publishers.

Allen, A. L., & Rothman, J. E. (forthcoming). Postmortem privacy. *Michigan Law Review*, *123*. Retrieved January 16, 2025, from https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4834871

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, 610–623.

Blaauw, M. (2013). The epistemic account of privacy. *Episteme*, *10*(2), 167–177.

Black, S. K. (2002). Encryption. In S. K. Black (Ed.), *Telecommunications law in the internet age* (pp. 327–387). Morgan Kaufmann.

Brake, E. (2023). How does stalking wrong the victim? *Ethics*, *134*(1), 4–31.

Chandler, C. (2025). Inside the black box of predictive travel surveillance. *Wired*. Retrieved January 25, 2025, from https://www.wired.com/story/inside-the-black-box-of-predictive-travel-surveillance/

Cohen, L. (2025). Online behavioral ads fuel the surveillance industry—here's how. *Electronic Frontier Foundation*. Retrieved January 24, 2025, from https://www.eff.org/deeplinks/2025/01/online-behavioral-ads-fuel-surveillance-industry-heres-how

Corbyn, Z. (2024). The AI tools that might stop you getting hired. *The Guardian*. Retrieved January 19, 2025, from https://www.theguardian.com/technology/2024/feb/03/ai-artificial-intelligence-tools-hiring-jobs

Cox, J. (2024). Hackers steal text and call records of 'nearly all' AT&T customers. *404 Media*. Retrieved January 28, 2025, from https://www.404media.co/hackers-steal-text-and-call-records-of-nearly-all-at-t-customers/

Cox, J. (2025a). Candy crush, Tinder, MyFitnesspal: See the thousands of apps hijacked to spy on your location. *404 Media*. Retrieved January 19, 2025, from https://www.404media.co/candy-crush-tinder-myfitnesspal-see-the-thousands-of-apps-hijacked-to-spy-on-your-location/

Cox, J. (2025b). Hackers mined AT&T breach for data on Trump's family, Kamala Harris. *404 Media*. Retrieved January 28, 2025, from https://www.404media.co/hackers-mined-at-t-breach-for-data-on-trumps-family-kamala-harris/

De Choudhury, M., Gamon, M., Counts, S., & Horvitz, E. (2021). Predicting depression via social media. *Proceedings of the International AAAI Conference on Web and Social Media*, *7*(1), 128–137.

Dretske, F. I. (1981). *Knowledge and the flow of information*. MIT Press.

Duhigg, C. (2012). How companies learn your secrets. *The New York Times Magazine*. Retrieved January 19, 2025, from https://www.nytimes.com/2012/02/19/magazine/shopping-habits.html?smid=url-share

Dwork, C. (2006). Differential privacy. In M. Bugliesi, B. Preneel, V. Sassone, & I. Wegener (Eds.), *Automata, languages and programming* (pp. 1–12). Springer Berlin Heidelberg.

Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), OJ L 119, 4.5.2016, p. 1–88 (2016, May 4). Retrieved January 19, 2025, from https://data.europa.eu/eli/reg/2016/679/oj

Fallis, D. (2013). Privacy and lack of knowledge. *Episteme*, *10*(2), 153–166.

Fried, C. (1968). Privacy. *The Yale Law Journal*, *77*(3), 475–493.

Friedman, J. (2017). Inquiry and belief. *Noûs*, *53*(2), 296–315.

Garoogian, R. (1991). Librarian/patron confidentiality: An ethical challenge. *Library Trends*, *40*(2), 216–233.

Gavison, R. E. (1980). Privacy and the limits of law. *Yale Law Journal*, *89*(3), 421–471.

Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, *23*(6), 121–123.

Goldman, A. I. (1967). A causal theory of knowing. *Journal of Philosophy*, *64*(12), 357–372.

Goldman, A. I. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge: New studies in epistemology* (pp. 1–25). D. Reidel.

Goldman, A. I. (1999). *Knowledge in a social world*. Oxford University Press.

Goldman, A. I. (2002). Reply to commentators. *Philosophy and Phenomenological Research*, *64*(1), 215–227.

GOV.UK. (2024). Horizon scandal factsheet: Post office (horizon system) offences bill. *Department for Business & Trade, Gov.UK*. Retrieved January 25, 2025, from https://www.gov.uk/government/publications/post-office-horizon-system-offences-bill-supporting-documents/horizon-scandal-factsheet-post-office-horizon-system-offences-bill

Hawke, P., & Berto, F. (2021). Knowability relative to information. *Mind*, *130*(517), 1–33.

Heikkilä, M. (2022). Dutch scandal serves as a warning for Europe over risks of using algorithms. *Politico*. Retrieved January 25, 2025, from https://www.politico.eu/article/dutch-scandal-serves-as-a-warning-for-europe-over-risks-of-using-algorithms/

Hill, K. (2024). Facial recognition led to wrongful arrests. So Detroit is making changes. *The New York Times*. Retrieved January 19, 2025, from https://www.nytimes.com/2024/06/29/technology/detroit-facial-recognition-false-arrests.html

Ichikawa, J. J., & Steup, M. (2024). The Analysis of Knowledge. In E. N. Zalta & U. Nodelman (Eds.), *The Stanford encyclopedia of philosophy* (Fall 2024). Metaphysics Research Lab, Stanford University.

Inness, J. C. (1996). *Privacy, Intimacy, and Isolation*. Oxford University Press.

Johnson, J. L. (1989). Privacy and the judgment of others. *Journal of Value Inquiry*, *23*(2), 157–168.

Kappel, K. (2013). Epistemological dimensions of informational privacy. *Episteme*, *10*(2), 179–192.

Lamdan, S. S. (2013). Why library cards offer more privacy rights than proof of citizenship: Librarian ethics and freedom of information act requestor policies. *Government Information Quarterly*, *30*(2), 131–140.

Lanzing, M. (2016). The transparent self. *Ethics and Information Technology*, *18*, 9–16.

Lanzing, M. (2019). "Strongly recommended" Revisiting decisional privacy to judge hypernudging in self-tracking technologies. *Philosophy & Technology*, *32*, 549–568.

Lundgren, B. (2021). How we can make sense of control-based intuitions for limited access-conceptions of the right to privacy. *Journal of Ethics and Social Philosophy*, *20*(3).

MacAskill, E. (2015). The NSA's bulk metadata collection authority just expired. What now? *The Guardian*. Retrieved January 19, 2025, from https://www.theguardian.com/us-news/2015/nov/28/nsa-bulk-metadata-collection-expires-usa-freedom-act

Mackie, J. L. (1977). *Ethics: Inventing right and wrong*. Penguin Books.

Macnish, K. (2018). Government surveillance and why defining privacy matters in a post-Snowden world. *Journal of Applied Philosophy*, *35*(2), 417–432.

Mainz, J., & Uhrenfeldt, R. (2020). Too much info: Data surveillance and reasons to favor the control account of the right to privacy. *Res Publica*, *27*(2), 287–302.

Marmor, A. (2015). What is the right to privacy? *Philosophy & Public Affairs*, *43*(1), 3–26.

Marmor, A. (2023). Privacy in Social Media. In *Oxford Handbook of Digital Ethics*. Oxford University Press.

Martin, N. (2018). How DNA companies like Ancestry and 23andMe are using your genetic data. *Forbes*. Retrieved January 29, 2025, from https://www.forbes.com/sites/nicolemartin1/2018/12/05/how-dna-companies-like-ancestry-and-23andme-are-using-your-genetic-data/

Matheson, D. (2007). Unknowableness and informational privacy. *Journal of Philosophical Research*, *32*, 251–267.

McDonald, A. M., & Cranor, L. F. (2008). The cost of reading privacy policies. *A Journal of Law and Policy for the Information Society*, *4*, 543.

Minde, T. B. (2023). Generative AI does not run on thin air! *RISE*. Retrieved January 29, 2025, from https://www.ri.se/en/news/blog/generative-ai-does-not-run-on-thin-air

Moore, A. D. (2003). Privacy: Its meaning and value. *American Philosophical Quarterly*, *40*, 215–227.

Moore, A. D. (2010). *Privacy rights: Moral and legal foundations*. The Pennsylvania State University Press.

Munch, L. A. (forthcoming). Why the NSA didn't diminish your privacy but might have violated your right to privacy. *Analysis*. Retrieved January 16, 2025, from https://philarchive.org/rec/MUNWTN

Munch, L. A., & Mainz, J. (2023). To believe, or not to believe – that is not the (only) question: The hybrid view of privacy. *The Journal of Ethics*, *27*(3), 245–261.

Murray, T., Cheong, M., & Paterson, J. (2023). The flawed algorithm at the heart of robodebt. *Pursuit from University of Melbourne*. Retrieved January 25, 2025, from https://pursuit.unimelb.edu.au/articles/the-flawed-algorithm-at-the-heart-of-robodebt

Nagel, T. (1998). Concealment and exposure. *Philosophy and Public Affairs*, *27*(1), 3–30.

Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Standford University Press.

Parent, W. A. (1983). Recent work on the concept of privacy. *American Philosophical Quarterly*, *20*(4), 341–355.

Patterson, A. E. (2020). *The social epistemology of prejudice & privacy: Toward new frameworks* [Doctoral dissertation, Cornell University].

Pepper, A. (2020). Glass panels and peepholes: Nonhuman animals and the right to privacy. *Pacific Philosophical Quarterly*, *101*(4), 628–650.

Pettit, P. (1996). Freedom as antipower. *Ethics*, *106*(3), 576–604.

Privacy International. (2019). Why am I really seeing that ad? The answer might be Real Time Bidding (RTB). *Privacyinternational.org*. Retrieved January 24, 2025, from https://privacyinternational.org/explainer/2974/why-am-i-really-seeing-ad-answer-might-be-real-time-bidding-rtb

Purtill, J. (2024). 23andMe is on the verge of bankruptcy. it may be too late to delete your genetic data. *ABC Science*. Retrieved January 29, 2025, from https://www.abc.net.au/news/science/2024-10-17/23andme-genetic-data-privacy-bankrupt-dna-test-ancestry/104455816

Radford, C. (1966). Knowledge: By examples. *Analysis*, *27*(1), 1–11.

Raz, J. (1988). *The morality of freedom*. Oxford University Press.

Reiman, J. H. (1995). Driving to the panopticon: A philosophical exploration of the risks to privacy posed by the highway technology of the future. *Santa Clara High Technology Law Journal*, *11*(1), 27–44.

Rössler, B. (2005). *The value of privacy*. Polity Press.

Russell, B. (A. W. (1948). *Human knowledge: Its scope and limits*. Allen and Unwin.

Salas-Zárate, R., Alor-Hernández, G., Salas-Zárate, M. D. P., Paredes-Valverde, M. A., Bustos-López, M., & Sánchez-Cervantes, J. L. (2022). Detecting depression signs on social media: A systematic literature review. *Healthcare*, *10*(2).

Schaake, M. (2024). *The Tech Coup: How to Save Democracy from Silicon Valley*. Princeton University Press.

Schleifer, T., & Ngo, M. (2025). Inside Elon Musk's Plan for DOGE to Slash Government Costs. *The New York Times*. Retrieved January 25, 2025, from https://www.nytimes.com/2025/01/12/us/politics/elon-musk-doge-government-trump.html?searchResultPosition=11

Skopek, J. M. (2020). Untangling privacy: Losses versus violations. *Iowa law review*, *105*(5), 2169–2231.

Solove, D. J. (2002). Conceptualizing privacy. *California law review*, *90*(4), 1087–1155.

Solove, D. J. (2006). A taxonomy of privacy. *University of Pennsylvania Law Review*, *154*(3), 477–564.

Staab, R., Vero, M., Balunović, M., & Vechev, M. (2024). Beyond memorization: Violating privacy via inference with large language models. *The Twelfth International Conference on Learning Representations*.

Studer, T. (2011). Justification logic, inference tracking, and data privacy. *Logic and Logical Philosophy*, *20*(4), 297–306.

Susser, D., Roessler, B., & Nissenbaum, H. F. (2019). Technology, autonomy, and manipulation. *Internet Policy Review*, *8*(2).

Thomson, J. J. (1975). The right to privacy. *Philosophy and Public Affairs*, *4*(4), 295–314.

Varian, H. R. (2009). Economic aspects of personal privacy. In W. H. Lehr & L. M. Pupillo (Eds.), *Internet policy and economics: Challenges and perspectives*. Springer US.

Véliz, C. (2020). *Privacy is power: Why and how you should take back control of your data*. Penguin Random House.

Véliz, C. (2024). *The ethics of privacy and surveillance*. Oxford University Press.

Warren, S. D., & Brandeis, L. D. (1890). The right to privacy. *Harvard Law Review*, *4*(5), 193–220.

Westin, A. F. (1967). *Privacy and freedom*. Atheneum.

Williams, A. (2024). Zombie trainers and a new era of forced labor. *DAIR*. Retrieved January 29, 2025, from https://data-workers.org/adrienne/

Zagzebski, L. (1994). The inescapability of Gettier problems. *The Philosophical Quarterly*, *44*(174), 65–73.

Zangwill, N. (2020). Epistemic pluralism. *Metaphilosophy*, *51*(4), 485–498.