On Privacy

Carissa Véliz
Faculty of Philosophy
Christ Church

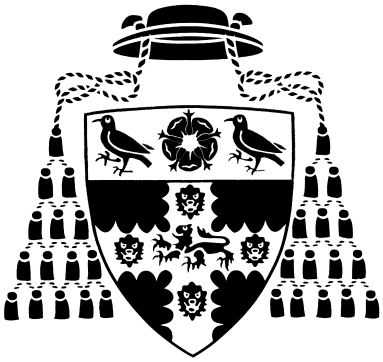

Supervisors:
Prof. Roger Crisp
Prof. Cécile Fabre

Thesis submitted to the University of Oxford
for the degree of DPhil in Philosophy
Trinity Term, 2017

To Sir, with Love

On Privacy

Carissa Véliz
Faculty of Philosophy
Christ Church
DPhil in Philosophy
Trinity Term, 2017

Abstract

This thesis concerns the ethics and political philosophy surrounding privacy. It investigates what privacy is, what is at stake in its loss, and how it relates to other rights and values.

The first part sets the groundwork for the rest of the thesis. Chapter One delves into the origins of privacy. I argue that privacy is not a recent cultural product, but rather a need buried deep in our evolutionary and human history.

The second part of the thesis is dedicated to conceptual issues. Chapter Two clarifies the relation between privacy and the public and private divide. I argue against the popular belief that privacy is an issue that belongs solely to the private sphere. Chapter Three reviews the most influential definitions of privacy that have been offered in the legal and philosophical literature, and points out some of their shortcomings and strengths. In Chapter Four, I develop my own definition of privacy as remaining personally unaccessed, as well as an account of the right to privacy as a right to a robustly demanding good. I also map out the moral significance of privacy perceptions, and privacy-related obligations.

The third part of the dissertation concerns practical issues. Chapter Five inquires into the relationship between security and privacy. I argue that mass surveillance is a disproportionate, unnecessary, and ineffective response to the threat of terrorism. I also argue that encryption should be widely used, as it can curtail the mass surveillance of content and protect people without seriously obstructing criminal investigations. Chapter Six explores the relationship between privacy and transparency. I argue that transparency should sometimes be limited in the interest of privacy. Chapter Seven deals with the questions of whether we can lose privacy to computer algorithms, and whether decision-making algorithms can violate our right to privacy. I answer both questions in the negative, as algorithms are currently neither our peers nor moral agents responsible for their actions.

The conclusion sketches some of the lessons learnt over the course of this investigation.

Acknowledgements

I have been the happy recipient of innumerable acts of kindness that have contributed to the completion of this dissertation. So many people have helped me along the way that the following list is bound to be incomplete. I hope that those who are missing will forgive me, and still accept my warmest gratitude for their generosity.

A graduate student could not ask for better supervisors than Roger Crisp and Cécile Fabre. They are both role models in their impeccable work ethic. Their feedback was, without exception, prompt and insightful, and their criticisms challenging and constructive. Working with them has been a genuine pleasure. It has proved invaluable to my philosophical training, and I am forever in their debt for their unwavering support.

I wish to acknowledge my College, Christ Church, the Faculty of Philosophy, the Uehiro Centre for Practical Ethics, and the Mexican Ministry of Education (SEP, DGRI) for their financial support.

I was fortunate to spend the last three years at the Uehiro Centre for Practical Ethics, amongst some of the most brilliant and kind people I have ever met. My deepest thanks to Julian Savulescu, without whose support, advice, and encouragement academic life would have been much more daunting. My thanks to Areti Theofilopoulou, for being the perfect officemate, and a friend I hope to keep for life. I am grateful to Josh Shepherd, Hannah Maslen, Jonny Pugh, Tom Douglas, Katrien Devolder, Neil Levy, Dominic Wilkinson, Guy Kahane, Will Davies, Antonio Diéguez, Francesca Minerva, Alberto Giubilini, Rebecca Roache, Chris Gyngell, and Gulzaar Barn, among others, for inspiring me through their ideas, and making work feel like leisure. Special thanks to all the administration team—Miriam Wood, Rocci Wilkinson, Deborah Sheehan, and Rachel Gaminiratne—who keep the world going round, and who made my life considerably easier and more pleasurable. Iris Geens and Sally Baume, from the Faculty of Philosophy, have also been tremendously helpful in their admin roles; many thanks to them.

In New York, I am indebted to Jesse Prinz, Carol Gould, and Graham Priest, for their fabulous courses and their support.

Parts of this thesis were presented at the Nuffield Workshop in Political Theory, the Surveillance Ethics Workshop, the St. Cross Special Ethics Seminar, the Applied Ethics Discussion Group, the Ethics Seminar at Ethox Centre, and the Research Seminar in Political Theory, all at the University of Oxford. Some chapters were also presented at the Medicine, Society, and Value Conference at Duke University, the MANCEPT Conference in Political Theory at the

University of Manchester, and the Workshop in Applied Ethics at the University of Bucharest. My thanks to the organisers and participants of those events.

My supervisors, mentors, and academic friends and colleagues, all tried their best to save me from the embarrassment of my mistakes. Any merit this thesis may have I owe to them. The faults, alas, can only be attributed to either my clumsiness, or my stubbornness in not following their advice.

This thesis is dedicated to Roger D. Gouran, my high school teacher, because he made it all possible. He found me at a moment in life when everything could have gone wrong, and cultivated in me the love of literature, film, and philosophy. His favourite movie was *To Sir, with Love*. With a teaching vocation the likes of which are the stuff of fiction, he took it upon himself to fill in the gaps in my education, starting with Shakespeare. I wish he could have lived to see the fruits of his efforts. His death is one of the greatest losses in my life.

I am grateful to all my friends and loved ones, living near and far, for being there for me at every turn. Thanks to them, I never lacked for warmth, understanding, or good fun. I am especially grateful to: Bent Flyvbjerg, for celebrating the completion of this thesis with me as if it had been your own, for seeing me and allowing me to see you, for sharing a passion for intellectual discovery and good writing (as well as bookmarks), for Magic Monday and beyond; Antonella Mallozzi, for being someone I can talk to about anything without a single euphemism, and for our New York adventures; Daniela Torres, for being my oldest and closest friend; David Ewert, for all your wisdom and kind-heartedness, and for staying close despite the long silences; David Rodríguez-Arias, for so many years of complicity and support, and for believing I could do anything I wanted with such steadfast conviction that you sometimes even managed to convince me; Diego Rubio, for our breakfasts, our dreams of better futures, and for reminding me, whenever I got caught up in trifles, that I had more important things to think about; Jorge Volpi, because the conversations we had while we walked the streets of New York continue to inspire me; Julia Powles, for your encouragement, wit, and a late summer night's dream; Luciano Espinosa, for making me appreciate philosophy at its best, for teaching me that I could find my homeland in books, and for more than a decade's worth of meaningful and heartfelt conversations, some of which have changed the course of my life; Marina López-Solà, for sharing your grandmother's wisdom with me, and being more important to me than you can imagine; Teresa López de la Vieja, for introducing me to ethics and political philosophy, for the many years of friendship, and the battles waged together.

I wish to express my gratitude to Aitor Blanco, Arancha San Ginés, Cecilia Tilli, Ilia Galán, James Williams, Kyo Ikeda, Lilian Bermejo and Javier Rodríguez Alcázar, Lorenzo Greco, all the organisers of the Madariaga Series, Pablo Serna, Rafo Mejía, Ricardo Parellada, Rosana Triviño, Jaime Alvar and Miguel, Theron Pummer, Txetxu Ausín, Nuria Ribes and Jaime, and Václav Janeček. Each of them supported me at some crucial moment or another in the past few years. For bringing magic and glee to my Oxford nightlife, thanks to Katja and Sergio, as well as Bent, Christian, Dario, Darío, Ellie, Elliot, Elvis, George, Gwyneth,

Hassaan, Ilias, Javi, Joachim, Julian, Kyo, Laura, Leo, Loic, Manuel, Nico (¡saltito!), Niels, Pablo, Paul, Phil, Philippe, Sissi, Solène, both Steves, and Zoltán. Three custodians at Christ Church made me smile every time I walked through the quad and never failed to ask about my wellbeing: thanks to Samuel, Kirk, and Clarindo.

Lastly, but most importantly, I owe more than I can express to my family. Gracias infinitas. A mi madre, María Perales, por darme tanto que ni un sinfín de palabras de agradecimiento podrían hacerle justicia a su generosidad; por su apoyo incesante, por impulsarme a volar, por ponerme siempre por delante de sus necesidades, por acompañarme al otro lado del teléfono cada vez que he llamado desde que me fui de casa, por ser mi compañera de viajes, por entender mis pasiones y perdonar mis ausencias. A mi padre, Héctor Véliz, por inculcarme el amor por la lectura, por ser la única persona que ha leído absolutamente todo lo que he escrito en filosofía, y por todos los libros que sin falta me regala cada año. A mi hermano Ivano, por sus abrazos y sus palabras, siempre afectuosas, siempre alentadoras; por ser un ejemplo de genialidad artística y humildad a la vez, por enseñarme que nadie vale más ni menos que otro, y que hay que apreciar y tratar con igual respeto al pez chico que al grande. A mi hermano July, por enseñarme desde pequeña que jugando se vive mejor; por estar siempre al pendiente, siempre disponible, siempre tan sólo a una jugada de distancia; por ser tan valiente y cuidar tanto de todos. A Ale, por haberse convertido en una hermana a través de los años, por los innumerables detalles que tiene conmigo, por querer tanto y tan bien a July. A Alexis, por su cariño y apoyo, y por cuidarnos a Ivano. Todos juntos me proporcionan un colchón en la vida que me ha permitido perseguir mis sueños por las alturas del mundo con la tranquilidad de saber que si me resbalo, caigo en blandito. No se puede pedir más de una familia.

Gracias a mi tío Remy, por salvarnos la vida más de una vez. Y gracias a mi familia valenciana—en especial a Silvia Gandía, Marisol Gandía, Sole, Maite y Rocío Vidal, Jesús Orera, y Chimo Perales—por acogerme. Finalmente, gracias a mis abuelos, José Perales y Carmen López, y a mi tío abuelo, Óscar López, por hacer posible que yo pudiera vivir las ilusiones que una guerra les negó.

Table of Contents

Introduction

In June 2013, Edward Snowden, then an NSA (National Security Agency) contractor, shocked the world by blowing the whistle on the existence of an extensive network of mass surveillance. With more than a million stolen documents to back up his claims, Snowden showed that the NSA was carrying out widespread electronic eavesdropping beyond American borders with the help of intelligence agencies in other countries (notably, the Government Communications Headquarters, or GCHQ, in the United Kingdom), and with the complicity of Internet and telecommunications corporations. Snowden revealed that the world was being spied on, and that the snooping covered data about telephone calls, emails, messages, location, financial transactions, browsing history, and more (Snowden 2014).

The Snowden revelations caught philosophers off guard. The philosophical literature on privacy is not extensive, and only a small proportion of it is contemporary enough to reflect post-Internet realities. Since 2013, privacy has hardly left the newspapers' headlines. It is one of the most pressing issues of our time. With the development of new privacy-invasive technologies such as drones, wearables, and the Internet of Things (everyday objects connected to the Internet and turning into smart devices); with a world economy that is increasingly fuelled by personal data; and with governments passing laws that continue to transform

the privacy landscape, philosophical reflection on the ethics and politics of privacy has never been more necessary.

In recent years, more philosophers have become interested in the topic, and gaps in the literature seem to be filling slowly. This thesis is intended to contribute to a better understanding of privacy—what it is, what is at stake in its loss, and how it relates to other rights and values. As I write these words, laws that affect privacy are being proposed in numerous countries; corporations, think tanks, and NGOs, among others, are developing ethical codes regulating the handling of personal information; the use of Big Data is growing in research, marketing, and even politics. The decisions we make about privacy today and in the coming years will shape the history of humanity for decades to come. Societal choices about privacy will influence how political campaigns are run, how corporations earn their keep, the power that governments may wield, the advancement of medicine, and the risks we are exposed to (e.g., identity theft, revenge porn, etc.). The least we ought to do is think carefully about these issues.

Structure

Part I: Groundwork

Chapter One explores the evolutionary, historical, and cultural origins of privacy. In this chapter I suggest that the human need for privacy is analogous to animals' need for territoriality. In the case of humans, however, and as a result of language,

a metaphorical zone of personal information complements the purely physical zone of personal space. I cite some of the anthropological observations that have been carried out on widely diverse cultures that suggest that privacy is a universal need among humans, even if what is understood by it may vary significantly from culture to culture. I also offer a brief overview of how privacy has been recognised and construed throughout history.

It may strike some readers as odd to find in a philosophy dissertation a chapter that is more historical and anthropological than philosophical. Empirical facts, however, matter a great deal in ethics, practical ethics, and political philosophy. In order to make normative recommendations, one must first have knowledge of the current state of affairs. Additionally, philosophy tests definitions of concepts against people's intuitions. Those intuitions arise in part out of the history of our culture and practices, such that, without historical knowledge, there is more danger of being misled. Furthermore, privacy is not (yet) a popular topic in philosophy, and background knowledge cannot be assumed. Finally, many of the misconceptions surrounding privacy that contribute to people not giving it the importance that, I will argue, it deserves, are sustained through ignorance of historical and anthropological facts. The usefulness of the chapter will become apparent in the rest of the dissertation as I cross-refer back to it.

Part II: Conceptual issues

Chapter Two is dedicated to dispelling a popular misapprehension: the belief that privacy is always related to the private sphere, and never to the public sphere. This common-sense view leads to unfortunate implications, such as the opinion that there should be no privacy protection on the streets. I argue that appealing to the private/public divide is unhelpful both in defining what privacy is and in determining what ought to be protected by privacy.

Chapter Three turns away from popular beliefs about privacy towards academic accounts of the concept. I review some of the definitions and accounts of privacy that are most influential in law and philosophy—those that turn on the concepts of the right to be left alone, control, possession of information, and limited access, as well as the attempt to understand privacy as a 'family resemblance' concept and the turn to reductionism. For each account, I point out its strengths and weaknesses. This chapter serves as preliminary work for the next one, where I present my own view of privacy. My account borrows from some of the virtues of the definitions presented here and tries to avoid their pitfalls.

Chapter Four offers a map of the conceptual and moral territory of privacy. I argue that privacy is best understood as not being personally accessed, and distinguish privacy from the right to privacy, which I argue, is a right to a robustly demanding good (i.e., one that has counterfactual demands). The chapter also

explores the value of privacy, the role of social norms, the moral significance of privacy perceptions, and the importance of privacy-related obligations.

Part III: Practical issues

Having explored conceptual issues raised by privacy and the right to privacy, in Part III, I bring my account to bear on three practical questions.

Chapter Five focuses on the relationship between privacy and security in the context of terrorist threats and mass surveillance. I first weigh security against privacy in an effort to assess the importance of the interests at stake. I then investigate whether mass surveillance is a proportionate, necessary, and effective measure in response to terrorism and argue that it is not. I then explore the internal connections between privacy and security and argue that these rights are less in conflict than is usually thought. The chapter ends with some reflections on the implications for encryption; I argue that encryption is not a significant obstacle to criminal investigations and should therefore be used widely.

Chapter Six explores the relationship between privacy and transparency. It delves into the pros and cons of radical transparency—the idea that every institution and individual should be able to surveil any other institution or individual at any time. I first present the supposed virtues of transparency and offer a picture of what a radically transparent society might look like. I then present some limits to the virtues of transparency and possible problems a radically transparent society

would face. From the perspective of privacy, a radically transparent society would be better than alternatives if either a) it is easier to protect privacy in such a society, or b) it has so many other advantages, that the benefits outweigh the sacrifice in privacy. I argue that neither condition obtains.

Chapter Seven deals with the questions of whether we can lose privacy to computer algorithms, and whether decision-making algorithms can violate our right to privacy. I answer both questions in the negative, as algorithms are neither our peers nor moral agents responsible for their actions. This does not mean, however, that algorithms are free from moral problems. The chapter then delves into some of these problems—in particular, the problem of opacity. When algorithms become too complex, not even their programmers can understand the decision-making processes they use. I conclude by arguing that algorithms are relevant to privacy because they are powerful tools that facilitate and enhance losses of privacy and violations of the right to privacy. Moral responsibility for those wrongs, however, still lies with humans, even when more and more of our personal information is being collected and analysed, not by humans, but by computer algorithms.

The thesis ends with a brief conclusion.

CHAPTER ONE

On the Origins and Nature of Privacy

If you are feeling tearful, embarrassed, unwell, or too exhausted to face the world, you probably feel like being at home, perhaps in your bedroom. Similarly, if you want to focus without distraction on a cognitively demanding task, you will probably seek solitude by closing the door to your office. When in the company of others you are not intimate with, you probably clothe your body and are inclined to keep certain kinds of personal information to yourself. Where do these tendencies come from? Are they recent? Are they restricted to our Western culture? This chapter intends to establish the groundwork for the rest of the thesis by exploring these questions.

In Section I, I explore the etymology of the word *privacy*. In Section II, I analyse the possibility that the human need for privacy may have its origins in propensities found in non-human animals. Finally, in Section III, I delve into the anthropology and history of privacy to explore variations and commonalities across time and cultures.

This chapter is meant, first, to provide a general sense of the meaning of privacy by exploring its etymology, and second, to argue that the desire for privacy is not

fundamentally a product of culture. Although culture does influence the way privacy is experienced and sought, the need for privacy arises from the kinds of creatures we are, much as the need for contact with our peers, and many other physical and psychological needs. I intend to demonstrate that this is the case by first showing how certain basic traits related to privacy can be found in non-human animals as well, suggesting that the need for some kind of privacy is universal to all animals and not only humans. I will then show how privacy needs break the boundaries of time and space in human communities, suggesting that privacy is not something recently invented by Western culture, as some seem to think. While these preliminary explorations into privacy do not establish its normative weight, they do suggest that privacy might not be something as dispensable as some critics suggest.

I. Etymology

Rigorous definitions of privacy will be explored in Chapters Three and Four. For now, an exploration of the etymology of the term will be enough to set the stage and give a rough and ready idea of what *privacy* means. Although current uses of the word *privacy* do not reflect all of the senses that it has carried in the past, we have inherited many of them, to varying degrees.

Privacy and *private* share the same Latin etymology. *Privatio* meant 'a taking away' (as in *deprivation*) and the adjective *privatus* correspondingly denoted something withdrawn from the public sphere (Webb 2007, xvi). Its oldest connotation,

therefore, was negative. The private was a lacking of the public—the public being the greatest good, where people fulfilled their potential. In this sense, a person who lived a strictly private life was not fully human (Aristotle 2013, I.2, Arendt 1998, 38).

In Cicero's vocabulary, to act *privatim* (the adverb whose opposite is *publice*) is to act not as a public officer, as a *magistratus* invested with a power originating from the people, but as a private individual. 'The private act was one committed not in the open, in the forum, before the eyes of all, but inside one's own house, in isolation, hidden from the view of others. The noun *privatum* refers to a person's own resources, property for his own use; and, again, to the home (*in private, ex privato:* inside or outside the home)' (Duby 1988b, 3-4). Thus, the fundamental meaning of *public* is that which the people as a community possess. In opposition, the *private* had the connotation of, on the one hand, that which departs from the communal and, on the other hand, that which relates to domesticity.

During the Middle Ages, the word *privatus* incorporated the meaning to be 'in retreat.' In a genealogy composed by Lambert of Saint-Omer in the early twelfth century, the term *privata* describes Robert le Frison's (count of Flanders) stay at the monastery of Saint-Bertin. This was a case of a *persona publica*, a prince, who temporarily abandoned his role as a sovereign to go into the relative seclusion of the monastery. Additionally, in the written Latin of the monasteries, the term *privatae* came to refer to latrines (Duby 1988b, 5-6). It was not so much that

latrines allowed one to be alone; rather, latrines were private insofar as they were away from public areas and behind closed doors.

In Romanesque language there seems to be a shift in the meaning of *private* that denoted the secret or the intimate. In the chronicle *Roman de Rou,* written by Wace in 1160-1170, Norman notables who are searching for ways to avoid taxes imposed on them by the Franks meet *privément*. Wace describes their gathering as one held behind closed doors. Thus the meaning of the word came to connote meetings done clandestinely. Privacy was that which enabled conspirators to plot against the ruling power. Private and public came to be understood as conflicting forces, and many times authorities saw themselves as having the duty to uncover private activities that could endanger their power (Duby 1988b, 6).

In English, the word *privacy* is a relatively recent term. The Oxford English Dictionary cites an isolated example of *privace* in the mid-fifteenth century. A similar term, *privity*, is an older word, used in the early thirteenth century. The likewise similar adjective and adverb *privy* and *privily* were used quite commonly in the fourteenth century (Webb 2007, xvi).

Ronald Huebert (1997) argues that the early modern meanings of the English terms *privacy* and *private* can be said to belong to four distinct semantic clusters. The first cluster understands what is private as a deficiency, a lack of what is public (hence the etymological cousins *privative, deprivation*). This sense of privacy

carries with it a negative connotation, as has been mentioned. In military terminology, for example, *private* designates the absence of rank.

A second semantic grouping is related to ownership and property. Writings from the sixteenth and seventeenth century consider private property as a refuge from the public sphere. Here we find a shift into a positive connotation of privacy.

A third cluster of meanings relates privacy to concealment of various kinds. *Privacy* can refer both to a concealed item and to the place of concealment. A special case is the modern phrase 'private parts.' In *A Relation of Some Yeares Travaile* (1634), Thomas Herbert writes that in some East Indian cultures 'the women goe most part naked, except a cloth which should cover those parts, made to be private' (cited by Huebert 1997, 32). He also makes reference to sexual activities being done in private.

The fourth semantic cluster that Huebert has identified has to do with various kinds of interiority. The poet Eliza speaks of the duty to share her creativity and not keep it isolated in her inner self: 'my desires were not given me, to be kept in private to my self, but for the good of others' (*Eliza's Babes or The Virgin's Offering,* 1652*)*. Huebert argues that, in Shakespeare's plays, the private refers both to being left alone and to being with (or in) one's thoughts. He offers the following quotation as an example: 'How is the King employ'd?' asks Suffolk in *Henry VIII* (2.2). 'I left him private, full of sad thoughts and troubles,' answers Lord

Chamberlain.[1] Often privacy is a reason for concern. Shakespeare seems to have had the sense that there is such a thing as too much privacy. In *Romeo and Juliet,* Montague is worried about his son, Romeo, who keeps too much to himself. According to Huebert, however, there is a shift in perceptions of privacy between the early sixteenth century and the mid-seventeenth century that goes from 'suspicion or hostility to privacy (…) to acceptance of and even a cherishing of privacy' (1997, 35).

While we seem to have lost the first semantic cluster, the other three senses are still found in our words 'private' and 'privacy.' In Chapter Three I will argue, however, that it is a mistake to confound private property and ownership with privacy, even if it is a common thought.

In nineteenth-century French dictionaries, the verb *priver* is found, meaning to tame or domesticate. Likewise, the adjective *privé* describes the family, the home, the domestic interior. In Émile Littré's dictionary, originally issued in 1863-72, one finds among other examples an expression coming into current use at the time: 'Private life should be lived behind walls.' Littré goes on to explain that '[i]t

[1] It seems to me that the quotations offered by Huebert as evidence for this fourth sense of privacy do not show incontrovertibly that they are referring to interiority. For example, Huebert quotes Shakespeare's Malvolio, in *Twelfth Night* (3.4), who says to Olivia: 'Go off, I discard you. Let me enjoy my private.' Huebert concedes that one interpretation of the quotation is simply that Malvolio wants to be alone, but, he adds: 'I'm confident that Malvolio's "private" here is also his inner self, his state of mind.' Regardless of whether Huebert is right that the modern meanings of privacy include a sense of interiority, this notion can be found alive and well in the present. For most people, there would likely be nothing more invasive of their privacy than another person who could read and divulge their thoughts.

[was] not permissible to inquire or talk about what goes on in the home of a private individual [*particulier*]' (Duby 1988b, 3).

With this general idea of what privacy has meant and means, I will next probe the possibility that the human desire for privacy has its origins in animal life.

II. Animal proto-privacy

Some scholars of privacy have suggested that human beings' need for privacy is likely rooted in our animal origins (Westin 1970, 8, Hirshleifer 1980, 655, 657)—that is, that it is not a purely cultural phenomenon, but rather stems from a more primitive source. In what follows I discuss four traits that are related to privacy and can be observed in humans and some non-human animals alike: the need to withdraw from others, the ability to deceive, the desire to save face, and the tendency to feel uncomfortable when others stare.

Withdrawal

With the exception of social insects, almost all animals seek individual seclusion or intimacy in small groups or couples at one time or another (Klopfer and Rubenstein 1977, 53). As I write this chapter, a video has been circling social networks showing a cat who likes to lock himself in a closet, shielding himself from

his flatmates (a dog and another cat).[2] An even more surprising example of animal withdrawal are babblers: even though these birds 'spend most of their lives within sight of their fellow group members, they do not copulate in their presence: they hide under or behind bushes and copulate in privacy' (Zahavi and Zahavi 1997, 145).[3] Chimpanzees have also been known to have secret mating *rendezvous*, but I will come back to these later.

Animal withdrawal from others often expresses itself as territoriality—when an animal claims a limited area for itself and defends it against intruders. Often the territory defended is an individual one, but it may also be the territory of a couple or a small group. Empirical evidence suggests that members of at least some territorial species use 'naturally occurring landmarks' to demarcate the limits of their territories (Eason, Cobbs, and Trinca 1999). Even though territoriality is widely spread across the animal kingdom, it may have different functions in different species, and it may not be possible to generalise a common function (Brown and Orians 1970, 248). In some cases, territoriality can be explained as an effort to secure food. In other cases, it is related to mating stations. In cases in which territorial behaviour has been shown to be unrelated to food, security, or mating, it is unknown why animals engage in it (Lawes and Henzi 1995, 242).

[2] See https://www.youtube.com/watch?v=_OG28uamCqA. Also see this compilation of cats seeking privacy: https://www.youtube.com/watch?v=XFCxqXiOprA. Accessed on December 29, 2016.

[3] I am grateful to Alex Kacelnik, Professor of Behavioural Ecology at the University of Oxford, for calling my attention to this example.

I intend this discussion of withdrawal by animals to count as support for the importance of privacy. It could be argued, however, that territorial behaviour did not necessarily evolve because it has a specific function to play (Brown and Orians 1970, 248), and that animals might be better off without it. Individuals acting territorially have been successful in reproducing, and that is how the feature has been preserved. Since natural selection selects for traits that promote survival, however, it is possible that a trait may be preserved that is detrimental to the wellbeing of the species, the individual, or both (Hirshleifer 1980, 653). There are traits that do not worsen the chances of survival for a species, but that a species might be better without. There are also traits that might have been helpful at some point in the history of a group or species but are no longer so. In short, natural selection does not select what is *ideal* for the individual or the species; it only has to select for what is *good enough* for survival and reproduction. Territoriality is a costly behaviour. It takes time and energy to defend one's space. One hypothesis is that animals would be better off without their territorial instincts. I think this hypothesis is wrong for the following reasons.

Peter H. Klopfer and Daniel I. Rubenstein (1977) propose that privacy be interpreted in economic terms across species. According to them, there is an equilibrium level of privacy that is reached when the marginal gain in fitness due to increased privacy (through territoriality, for example) equals the marginal loss of fitness due to the costs of increasing privacy (64). This interpretation presupposes there must be *some* gain in fitness due to increased privacy.

Having an adequate amount of personal space, which is closely related to territory and privacy, seems to be a requisite for the wellbeing of groups and individuals. Wellbeing may be one of the main gains derived from having a territory. Territories give individuals more chances for withdrawal. When animals get crowded, they become stressed. Hindering animals' desires for withdrawal has catastrophic effects. High population density creates high blood pressure, multiple diseases, aggressive behaviour, and in some cases may even cause death (Andrews 1979, Christian, Flyger, and Davis 1960). At least part of what is important about having personal space is having the chance to rest. Crowded broiler chickens, for example, try to remain close to walls in order to avoid disturbances from others as much as they can (Buijs et al. 2010). When conditions become too crowded, chickens become stressed and aggressive, pecking each other to death. To avoid such behaviour, a common practice is to sear their beaks (Ellis 2007).

Despite our cultural sophistication, humans are also animals—we share genes and evolutionary history with other animals. It is therefore likely that the human need for privacy is related to animal territoriality and the need for personal space, and that it likewise contributes significantly to our mental and physical wellbeing. In the case of humans, however, as a result of the development of language, a metaphorical zone of personal information complements the physical zone of personal space. For humans, privacy is not only a matter of periodic physical withdrawal, but also of keeping certain kinds of information to ourselves. There are two ways of keeping others from knowing certain things about us: withholding information (keeping quiet, hiding information) or deceiving (lying, misleading

others). Depending on the context and the reasonable expectations of the parties involved, withholding information may also be conceptualised as a form of deception. Can animals deceive as well?

Deception

Although it is a matter of controversy how much animals understand about the tricks they engage in, deception is a fundamental tactic to gain advantages in the animal world. In order to attract potential female mates, male chickens produce food-associated calls even when there is no food available (Gyger and Marler 1988). When pitted against dominant chimpanzees in contests over food, instead of being open and honest about their consumption of food, subordinate chimpanzees will selectively obtain pieces of food that dominant individuals had not seen or did not know about (Hare et al. 2000). When Rhesus monkeys are given the opportunity to steal a grape from a human who is looking away, they will choose the grape that is placed in a container that makes no sound when touched (as opposed to the one placed in a container that makes noise when tinkered with) so as to pass unnoticed (Santos, Nissen, and Ferrugia 2006). These studies suggest that some animals, and especially non-human primates, may have an understanding of others' mental states. While it would be too bold to claim that animals have personal information that they do not want to share with others, the seed of that phenomenon, withholding information and deceiving, does seem to be found among animals.

Whiten and Byrne (1988) have created a taxonomy of deception in primates that includes deceiving others through concealing something from them (withholding information), distracting them (which sometimes involves seduction), and faking of various types. Some of the examples discussed in their work do not seem very related to our common sense understanding of privacy (e.g., a female seducing a male and unexpectedly stealing his food and running away). Some other examples, however, would certainly be considered cases related to privacy if only the protagonists of the anecdotes were human, as in this example, taken from Frans de Waal's *Chimpanzee Politics*:

> Dandy and a female were courting each other surreptitiously. Dandy began to make advances to the female, whilst at the same time restlessly looking around to see if any of the other males were watching. Male chimpanzees start their advances by sitting with their legs wide apart revealing their erection. Precisely at the point when Dandy was exhibiting his sexual urge in this way, Luit, one of the older males, unexpectedly came round the corner. Dandy immediately dropped his hands over his penis concealing it from view. (de Waal 2000, 36-37)

Granted, a man would not court a woman in the same fashion (one would hope), but hiding sexual arousal is undoubtedly a common manifestation of the human desire for privacy. De Waal mentions other cases in which chimpanzees hide their erections. Similarly, Whiten and Byrne (1988, 236-237) recount how gelada baboons suppress their usually loud vocalisations when mating within auditory (but not visual) reach of the rest of the group.

The reasons behind such examples of modesty, however, are potentially very different in humans and non-human primates. Dandy, the chimpanzee, is most

likely trying to avoid a dangerous confrontation with an older and more powerful male. De Waal believes that secret chimpanzee *rendezvous* are meant to avoid interruptions and male competition. It is not clear that this is the *raison d'être* of privacy in human sexuality. One could argue, however, that the lover of an adulterous woman likewise desires privacy in order to avoid having a threatening confrontation with her husband (or his own wife, if he is also married). The political sociologist Barrington Moore speculates that it is important for people to have sexual acts be private, among other reasons, to avoid unpleasant interruptions, and to control jealousy and possessiveness, which can lead to aggression (1984, 70-71).

Even if it is implausible to say that chimpanzees care about privacy in the way that humans do, they do exhibit similar forms of behaviour. Thus, it is plausible to think that, in order to understand the origin of human behaviour regarding privacy, we ought to look at animal behaviour such as that of chimpanzees. In the human case, privacy-related practices have taken different forms in different cultures, but the origins of such actions may still be found in animals.

Above and beyond the possible social functions there might be for sex being private, most people would *feel* embarrassed if any part of their sexual lives were exposed. Can chimpanzees or other non-human primates feel embarrassed?

Saving face

Chimpanzees do not seem to feel generally embarrassed about their sexuality. Some biologists, however, do think primates are worried about saving face and can feel embarrassment in other situations.

Jane Goodall tells the story of Freud, a five and a half year old chimpanzee who was swaying back and forth on the stem of a wild plantain, showing off near his uncle, an alpha male named Figan. When the stem of the plantain suddenly broke, Freud tumbled into the grass, unhurt but embarrassed, immediately looking over to see whether Figan had noticed (Goodall 2000, 166-167).

Similarly, Marc Hauser recounts how a male rhesus monkey fell into a ditch as he walked away from a female with whom he had just mated. The male quickly stood up and looked around, as if worried that the female had seen him. After making sure his fall had gone unnoticed, he resumed his walk, with a proud bearing (Hauser 2000, 200-201).[4]

What these and other cases seem to suggest is that chimpanzees (and perhaps other primates) worry about what others think about them and may experience embarrassment. Most biologists, however, seem to think embarrassment in non-human animals is not something we can detect, except by questionable subjective

[4] I include Hauser's example on the assumption that most of his work is accurate. It should be acknowledged, however, that in 2010, Harvard University found him guilty of fabricating data and manipulating the results of some of his studies.

judgments of similarity with human behaviour.[5] It thus remains unclear whether non-human animals experience embarrassment as we do. What is clear, however, is that others' gaze can make non-human animals uncomfortable, if not embarrassed.

Uncomfortable stares

If you were born somewhere in the West, at some point one of your parents or guardians probably taught you that one must not stare at people. It is not polite. Other people's gaze can make us feel self-conscious, nervous, and awkward. Among at least some primates, a steady, direct gaze is a component of threatening behaviour.

Primates such as gorillas interpret direct stares as challenges. If you ever go into a laboratory that experiments on non-human primates, you will be told to not look them in the eye, as this will upset them. After an unfortunate incident in 2007 with a gorilla named Bokito at a zoo in Rotterdam, a local health insurance company distributed more than 2,000 *BokitoKijkers* ('Bokito viewers'): cardboard glasses with holes in the middle and eyes printed on them that are looking upwards (Wang 2009). The glasses allow people to look straight at gorillas without them noticing the direct stare.

[5] Alex Kacelnik, Professor of Behavioural Ecology at the University of Oxford. Personal communication.

Predators gaze upon their prey before attacking, and it may be that the human discomfort at being looked upon comes from this fact in our evolutionary history. The less one is looked at, the less likely it is that one may become a prey. It might be thought that, if this is right, the desire not to be looked at is simply a remnant of an evolutionarily adaptive feature that is no longer relevant in today's world, since most people are no longer vulnerable to becoming prey to animal predators. It is still the case, however, that others who observe us make us more vulnerable to them. Consider how terrifying it can be for a woman to be observed lustfully by a man in a dark and deserted alley. The more someone watches us, the more information they can glean to use against us. Humans may not be prey to lions any more, for the most part, but we can still be victims of each other. Even in cases in which we are certain we are safe, as a result of our evolutionary history, we might not be able to shake off the feeling of discomfort and vulnerability we get when others look at us. The unease may also derive from our caring about what others think about us, so that we may act differently in their presence. I will come back to this point in Chapter Four.

Conclusion

While the observational evidence cited in support of the hypothesis that the origin of the human desire for privacy may be found in animal tendencies is certainly suggestive, claims in this regard can only be tentative. There is no easy way to dismiss worries regarding the possible anthropomorphisation of animal behaviour. That some acts by animals are similar to human acts motivated by the desire for

privacy does not mean that the underlying origins, functions, and mechanisms are alike (Klopfer and Rubenstein 1977, 53).

That being said, animals' needs for withdrawal and personal space do seem extremely similar to our own needs. The negative effects of crowding in prisons, for example, suggest that lack of opportunities for withdrawal are just as detrimental to human health and wellbeing (McCain, Cox, and Paulus 1976, García-Guerrero and Marco 2012). It seems, therefore that just like other animals, we need some space away from others simply in virtue of the kinds of creatures we are.

The case of deception is less clear. While we can evidently see our own abilities to withhold information and misguide others reflected in non-human animals, it is doubtful that they use these abilities to protect their privacy in the way we do. Other primates may share with us some abilities to fake and deceive—abilities that help us protect our privacy—but our use of those abilities is much more extensive. The most that can be asserted is that our evolutionary history has provided us with abilities to deceive that aid us in protecting our privacy.

In the case of embarrassment, anecdotal evidence suggests that non-human primates may exhibit similar emotions, but conclusions in this regard seem highly speculative and interpretative, so we have reason to be cautious.

Finally, the uneasiness experienced by primates when they are being stared at seems to be quite similar to the discomfort felt by humans. It is likely the case that direct gazes feel like a challenge or threat to most primates, and at least in the case of humans, the discomfort may also have to do with worrying about what others think about us.

That needs for privacy (or proto-privacy) can be found in non-human animals already suggests that the human desire for privacy is not a cultural product. However, this inference is not conclusive, as some of the evidence is anecdotal and subject to different interpretations. If privacy practices were to be found in different cultures and times, that would constitute further and much stronger evidence towards the universality of privacy.

III. Privacy: variations and commonalities across time and cultures

Sceptics of the universality of privacy are quick to point out that our ideas about privacy are recent, that at other times and cultures, privacy has not been valued like it is valued today. Even within what can be broadly construed as 'Western culture,' norms about what should remain private have undergone such dramatic changes over the course of history that the idea that privacy is a universal value or need is often challenged.

Lawrence M. Friedman, for example, believes that '[p]rivacy, as idea and reality, is the creation of a modern bourgeois society' (Friedman 2007, 258). Vinton Cerf,

one of the architects of the Internet and Google's Chief Internet Evangelist, thinks that privacy emerged during the industrial revolution (cited by Ferenstein 2013), and journalist Thomas McMullan (2015) has suggested privacy was invented with the telegraph. If privacy is indeed a recent cultural development, and people in the past and other cultures have lived happily in the absence of the kind of privacy that we now value, then it would seem more plausible that we could easily do away with it without much loss. In what follows, I go through some of the changes and contrasts that have taken place in Western societies. I then look at other, non-Western cultures. I argue that privacy is a desire that has persisted in diverse cultures and times.

Perhaps the most private of places for most contemporary Westerners is the bathroom. It is very rare for people to want company when they are sitting on the toilet. One of the historical examples often given to show that privacy is a recent development is Roman bathrooms (e.g., Ferenstein 2013). Romans were fond of combining 'evacuation and conversation,' as Bryson puts it (2010, 500). Their public latrines could seat twenty or more people close together. This custom lasted for centuries, with some modifications. Romans were not the only ones who were not shy in the bathroom. Hampton Court Palace, built in 1514, contained a 'Great House of Ease', which accommodated up to fourteen people; Charles II was accustomed to go to the 'privy' with two attendants; and George Washington's home had a lavatory with two conjoining seats (Bryson 2010, 500).

One might be tempted to infer that perhaps Romans (and people who followed their customs) did not have a need for privacy. If they were willing to do in public what today we view as one of the most private activities there are, then perhaps Romans did not feel the need to keep anything at all private. Maybe they did not feel embarrassment like we do, and were happy to share everything with the people surrounding them. There is much evidence against that hypothesis, however.

It is quite clear that the Romans were intent in keeping some kinds of information private. They frequently attended temples to ask favours from gods. Some of those favours were less virtuous than others. When they asked to be richer than their neighbours, for example, '[t]hey did not dare to utter such a wish out loud, in front of other worshipers, so they wrote it down and left the sealed document on the altar' (Veyne 1987, 211). Similarly, privacy of correspondence was highly valued by the Romans. When Mark Antony read out loud the letters that Cicero had sent him, the latter went before the Senate to denounce him:

> But he also read letters which he said that I had sent to him, like a man devoid of humanity and ignorant of the common usages of life. For who ever, who was even but slightly acquainted with the habits of polite men, produced in an assembly and openly read letters which had been sent to him by a friend, just because some quarrel had arisen between them? Is not this destroying all companionship in life, destroying the means by which absent friends converse together? How many jests are frequently put in letters, which, if they were produced in public, would appear stupid! How many serious opinions, which, for all that, ought not to be published! Let this be a proof of your utter ignorance of courtesy. (Cicero 2009, 18)

Critics might be tempted to think that, even if Romans valued privacy with regards to certain kinds of information or documents, they might not have had a desire for physical privacy. In other words, maybe ancient Romans did not have a need for physical withdrawal. While it is quite likely, if not certain, that Romans were more gregarious than modern day Westerners, there is evidence that suggests there were places and times for solitude. Veyne points out that '[f]or the ancients, a man's study was a sanctuary of private life' (1987, 229). The study was a place where a person could retreat from social life in order to read and write.

Perhaps, then, it was only privacy with regards to the body and bodily functions that the Romans did not value. After all, many times men and women bathed together in public baths. Considering their design of latrines and their bathing customs, it is plausible to think that Romans had no reserve when it came to exposing the body. Even this, more moderate hypothesis, is highly questionable. It is still a matter of debate to what extent people using public baths were unclothed (Fagan 2005, 25-26). It seems that there was much variety in this and other aspects of bathing. At least in some cases, there were different times or separate bathing wings for women and men. Furthermore, there is evidence of some prevalent taboos when it came to sex and exposing the body. A man was considered a libertine if he made love before nightfall,[6] without darkening the room first, and if he made love to a woman who was completely naked. According to Veyne, 'only fallen women made love without their brassieres, and paintings in

[6] Romans probably inherited this taboo from the Greeks. Moore points out that '[a]ccording to Pindar, both gods and men should have the modesty and restraint not to consummate a marriage in the light of day' (1984, 140).

Pompeii's bordellos showed even prostitutes wearing this ultimate veil' (1987, 203).

One might think that I chose the wrong society. If one is looking for a Western culture in which people do not have reserves about exposing their bodies, perhaps one should look to ancient Greece. The Greeks were well known for adopting in the seventh century B.C. the curious custom of practising athletics unclothed (Scanlon 2002, 326). Athletic nudity was unique to Greeks at the time. It differentiated them from other Mediterranean societies, and other nations mocked them for it (Scanlon 2002, 208). Despite this custom, however, ancient Greeks were vulnerable to being embarrassed about their naked bodies. It was not easy for people to accept the new custom of engaging in athletics in the nude, and it seems to have attracted some ridicule at the beginning from unconvinced Greeks (Scanlon 2002, 207-208). Furthermore, the prospect of disrobing in front of people, even in the context of an accepted practice, could still be a source of anxiety. First, Greeks worried about being physically fit. Lucian, for example, mentions how 'expecting to appear unclothed before so many people, [Greek athletes] try to attain good physical condition so that they may not be ashamed of themselves when they are stripped' (cited by Scanlon 2002, 208). Second, they worried about sexual arousal. It seems that the practice of infibulation (drawing up the foreskin of the penis and tying it with a cord) was widespread. The purpose was apparently 'to prevent the embarrassment of an erection in public context' (Scanlon 2002, 235). It is untrue, then, that the Greeks had no reservations about exposing their bodies.

Ancient Greeks were also concerned about private information. In *To Demonicus*, a text that offers advice to a young man, Isocrates encourages self-restraint. He advises one not to 'expose' oneself to others. Isocrates instructs Demonicus that if he ever needs advice from a friend about an embarrassing matter, he should pretend to be asking advice on behalf of someone else in order to get the benefit of the counsel without having to expose himself. 'Hence the Athenian private gentleman,' concludes Moore, 'was expected to be a very reserved and private person' (1984, 132). Not surprisingly, it was considered a disgrace to have domestic disputes amongst relatives become public (Moore 1984, 139).

What the Greek and Roman examples show is that, despite the widely different practices of our ancestors, we find at least as many commonalities as differences between their concerns about privacy and our own. The desire for privacy is not as new as some critics would like to make it seem. Ancient Greeks and Romans were familiar with feelings of embarrassment concerning the body, they valued privacy with respect to certain kinds of information, and they withdrew in solitude (or intimacy) to rest and to engage in activities such as writing. Similar examples can be cited throughout the history of the Western world.

Architecture is a telling expression of concerns about privacy. In their deliberate design of space, cultures reveal their priorities. At the end of the Middle Ages and the beginning of the Renaissance in Italy, when some people, the bourgeoisie, began to have enough money to build houses to their liking, they built private

bedrooms for all members of the family, including children. Sometimes even couples had different bedrooms for each spouse. The famous architect Leon Battista Alberti

> recommended that husband and wife each have a room so that neither would burden the other unduly—in case of illness, oppressive heat, or pregnancy, for example. The two rooms should communicate so that husband and wife could meet without attracting the attention of gossips. A quiet, heated, private room was even more indispensable to an elderly person (…). Most of all, however, a private room was needed by the head of household, particularly if he belonged to a great lineage. The bedroom was the secret chamber, where the master of the house contemplated his most precious possessions and consulted his most valued family documents as he decided on a proper course of action. (cited by de La Roncière 1988, 216-217)

It can be argued that the progression towards spatial privacy has been an on-going tendency on a par with material development. In 2011, about a third of households in the United States and the United Kingdom had one person living in them. In Sweden, where there is a strong welfare state that makes people not have to depend directly on one another, 47% of households have one resident (Klinenberg 2012). It seems that, in so far as people do not depend on others and can economically afford it, they choose ways of living conducive to privacy.

By citing these examples, however, I do not mean to give the impression that privacy has always been unquestionably valued in Western history. It is worth keeping in mind that the etymological origins of the word have negative connotations. At different times in history and different contexts, privacy has been regarded as something suspect and dangerous. In late antiquity, for example, Jews and early Christians worried that privacy could jeopardize loyalty and solidarity

with the group. The hope was that individuals would give the whole of themselves to their religion:

> In the first century A.D. this model [of suspicion about privacy] was supported, with widely varying degrees of urgency and abruptness, by the belief that through the action of God a social state presently governed by the abrasive opacities of double-heartedness would give way, among a true remnant of Israel, to a time of utter transparency to each other and to God. In such a true, redeemed community, the tensions of the "evil heart" would have been eliminated. (Brown 1987, 254-255)

Similarly, the familiar view that good people have nothing to hide and therefore only evil people want privacy can be found already in eighteenth century literature (Spacks 2003, 49).

Sometimes negative views of privacy have been associated specifically with women. Men in feudal France thought women were up to no good when they were alone in the *chambre des dames*, where it was accustomed to lock them up (Duby 1988a, 77-79). During the eighteenth century women were seen as vulnerable creatures who were endangered by privacy. It was thought that if they were allowed to read by themselves, they might be victims to uncontrolled fantasies that could lead to disaster (Spacks 2003, 10).[7]

[7] That men had more opportunities for privacy than women is an instance of a more general tendency of privacy to be a privilege of the more advantaged members of society. The rich—who can afford to have more personal space—typically have enjoyed more privacy than the poor. A critic might point out that, historically, rich people have had less privacy than others because they have been surrounded by slaves or servants. Slaves, however, did not count as equals (Veyne 1987, 73). Masters did not need to be worried about upholding social norms in their presence. In the words of Peter Brown, '[n]udity before one's slaves was as morally insignificant as nudity before animals' (1987, 246). As the lower echelons of society transformed from slaves to servants and their moral significance started

In short, privacy has not always been looked at favourably. But the mere fact that it has attracted criticism shows that there has been a persistent desire and practice in Western history to withdraw from others. If people did not seek privacy, critics would not have needed to express their worries. What examples—both of practices that protect privacy and condemnations of privacy—show is that privacy is not something we have invented recently. For all that has been said so far, however, it may still be the case that privacy is a Western invention, and not a panhuman desire. Can privacy concerns be found in societies that are culturally distant from our own?

In *Privacy: Studies in Social and Cultural History* (1984), Barrington Moore has convincingly answered this question in the affirmative. He understands privacy as 'a desire for socially approved protection against painful social obligations' (Moore 1984, 6). Moore looks at societies in which, at a first glance, privacy seems scarce or non-existent.

The Utkuhikhalingmiut (abbreviated as Utku), are an Eskimo community of twenty to thirty-five people who live inside the Artic Circle in Northern Canada. They are an isolated community, with the nearest trading settlement being several days away by dogsled. The Utku remind us that for much of human history,

---

approximating that of an equal, different privacy norms developed. In Victorian Britain servants were expected to keep out of sight and be as near invisible as possible (Bryson 2010, 147-148). Sometimes servants were made to turn their faces to the wall when their masters passed by. Not surprisingly, houses were increasingly designed to keep staff separate from family members.

solitude was very dangerous. Because of the hostile conditions they live in, the Utku cannot afford to be by themselves. They live in tightly packed tents (in the summer) or igloos (in the winter). Living in such close quarters, the Utku have devised an unusual form of withdrawal. In the igloo, every person has a sleeping space into which no one intrudes without permission. An individual seeking withdrawal may spend hours or even days lying silently in their bed, facing the wall, ignoring those around him, without being disturbed (Moore 1984, 4-11). When the environment does not allow for physical walls, social norms can act as invisible walls to protect privacy.

The most extreme culture explored by Moore is that of the Siriono Indians, a hunting and gathering people in Bolivia. The Siriono Indians are the likeliest candidates to have no concerns for privacy. A striking feature of this community is the utter disregard they show towards one another (Moore 1984, 18). They survive with a minimum of cooperation. In terms of privacy, *all* physiological activities may be carried out in public. According to Moore, 'Siriono society lacks all but the most minimal distinction between what is private and public—and suffers the disadvantages of both. There is neither protection against intrusion nor the advantages that come from cooperation and the recognition of a collective interest' (Moore 1984, 19). Even in this society, however, Moore finds embers of privacy. Couples often find seclusion in the forest to engage in sexual intercourse (Moore 1984, 67). Likewise, a desire for privacy can be glimpsed in the frequent complaints about noise and disturbances that are due to living in tight quarters (Moore 1984, 275). Finally, one can speculate that individuals welcome the break

they get from each other when they go out for a day of hunting or gathering, alone or in pairs. In many societies, hunting and gathering activities provide a temporary escape from the burdens of sociality and thus have a privacy function beyond providing food for the group.

To further support his observation that privacy is a panhuman desire, Moore cites the work of Clellan S. Ford and Frank A. Beach. They studied twenty-five societies about which adequate information about sexual habits existed. It turned out that twenty-four of those societies chose to have sexual intercourse in private places (Ford and Beach 1951, 70-71).

As a final non-Western example, it is worth mentioning that people in ancient China were not strangers to privacy concerns. Among their many norms, it was customary to consider the sexual act a sacred practice that should never be performed in public or talked about with strangers (Moore 1984, 256). In the *Analects*, there is a passage that suggests that ancient Chinese did not approve of intrusions into their private space. Confucius criticises a timid man who is pretending to be fierce because he is like a man who is so 'dishonest as to sneak into places where one has no right to be, by boring a hole or climbing through a gap' (Confucius 1938, 2:21).

After this brief tour of privacy through times and cultures, it seems warranted to believe that privacy is not a new concern. The desire to seek refuge from the demands and risks of sociality has been shared by many cultures throughout

history. This desire, of course, is complemented by an equally strong need to have company and receive comfort from our peers. Though both needs are shared throughout cultures and time, the specifics of how people attempt to satisfy both desires may vary widely. Although the desire for some forms of privacy can probably be found in most human beings across time and space, it is one that can be manipulated relatively flexibly through culture and other contextual conditions (e.g., when we are in absolute need of others for survival, privacy becomes a secondary worry).

Privacy is a need born out of sociality. As Moore points out, '[w]ithout society there would be no need for privacy' (1984, 73). We need others to survive and flourish, but given our nature, we also need to have some distance from others. For individuals to thrive, a balance must be struck between sociality and privacy, and the exact point and style of equilibrium varies both across societies and individuals. Ways of finding spaces and time for privacy also vary widely across cultures, but what remains common to all humans is the need to keep to ourselves from time to time and in certain circumstances.

In this chapter evidence has been presented against the view that privacy is a recent Western product, and therefore easily done away with. Another common myth associated with privacy is that all that is or should be private belongs to the private sphere, and that anything that goes on in the public sphere has no claim to privacy. I tackle this myth in the following chapter.

CHAPTER TWO

Privacy, the Public, and the Private

Before turning to the most influential definitions of privacy in the academic literature, I would like to clarify the relationship between privacy and the public/private spheres in order to avoid misunderstandings. When discussing privacy in social settings, I have often encountered people defending the view that there is no claim to privacy in the public sphere. The thought, then, is that, if something happens to someone while in the public sphere, or if some event or person belongs to the public sphere, they cannot be expected to enjoy any degree of privacy. I will argue that appealing to the private/public divide is unhelpful both in defining what privacy is and in determining what ought to be protected by privacy. The objective of this chapter is to show that it is not a contradiction in terms to say that privacy issues are not concerns that only and always belong to the private sphere.

In Section I, I consider William Parent's view as an example of a philosophical account that relies on the public/private divide to determine what privacy is. I analyse his proposal because it makes explicit ideas that are popular in Western countries but are often expressed with less clarity in the media or in conversation.

I go through some of the problems this kind of view encounters when defining privacy and justifying what ought to be private.

One might think that the problems with Parent's view stem from an incorrect way of dividing the public from the private. In Section II, I examine three alternative ways of marking out the divide—spatially, financially, and with respect to the roles we play. I argue that none of these ways of separating the private and the public can justify a definition or a defence of what ought to belong to the realm of privacy.[1]

In this chapter I will not yet offer a philosophically rigorous definition of privacy. As a rough and ready working understanding of privacy, it can be said that the private, in the sense of privacy, refers to that which is kept hidden by individuals from most other people (except a chosen few). I will mostly rely on paradigmatic privacy cases to argue that, whatever definition of privacy we favour, and as long as we want it to accord with common usage, we cannot rely on the private/public distinction as it is commonly understood in order to justify a claim about what privacy is and what ought to be kept private.

[1] Helen Nissenbaum argues for something similar, although she frames the private/public dichotomy in terms of incursions by government actors, incursions in private domains, and dissemination of private information (2010, 114-116). She believes technological advancement is the reason the dichotomy has become less useful as a foundation for a conception of privacy (2010, 116). I am not sure the dichotomy was ever accurate enough to be in alignment with social norms governing privacy.

I. Relying on the public/private divide: William Parent's account

William Parent has defined privacy as

> the condition of not having undocumented personal knowledge about one possessed by others. A person's privacy is diminished exactly to the degree that others possess this kind of knowledge about him. (Parent 1983a, 269)

For Parent, a fact is documented if it belongs to the public record (information published in newspapers, court proceedings, official documents open to public inspection, etc.). Nothing in the private sphere is therefore documented. Crucially for his definition of privacy, he understands personal information as

> facts which most persons in a given society choose not to reveal about themselves (except to close friends, family,…) or facts about which a particular individual is acutely sensitive and which he therefore does not choose to reveal about himself, even though most people don't care if these same facts are widely known about themselves. (Parent 1983a, 270)

Parent's account commits him to an extremely counterintuitive implication: that privacy is not lost through the publication of personal information. B loses his privacy the moment A gains some undocumented personal knowledge about him, but not when he puts that information in the public domain by publishing it, and not when other people start learning about the now documented personal information. For Parent, then, having one's affairs published on a tabloid is not a privacy problem.

One option to avoid this counterintuitive implication would be for Parent to defend the view that undocumented personal information should never become

documented or published. There is no evidence to think this is his view, and it would be a very unpalatable stance to take, as it would seriously diminish free speech and journalistic enterprises. If, however, he does not take this view, then the implication is that he is not able to claim, on the basis of his account of privacy, that some things which are part of the public domain should not be—things that, for privacy reasons, should have never been made public. Because, in his view, people only lose privacy when knowledge is first gained, Parent can make no assertions about the moment the information becomes part of the public record through documentation or publishing.

Despite these implications, Parent does say that 'a person does lose a measure of privacy at the time when personal information about him first becomes a part of the public record, since the information was until that time undocumented' (1983a, 271). I do not see how this can follow from his account, however. On Parent's definition, if person B only loses privacy when person A has *knowledge* of B's undocumented personal information, then B would lose privacy, for example, when a journalist *learns* something personal about him, not when he publishes that information. Gaining knowledge of B's personal information is not the same thing as that information becoming part of the public record. According to Parent's account, when people read about the information on the news, B does not further lose privacy because that information is part of the public domain the moment it is published.

The unfortunate implication of Parent's view is that privacy cannot be further lost once a piece of personal information is on the public record—even if it had been forgotten and is suddenly rediscovered. Suppose, says Parent, that A finds a story about B in an old newspaper in which he learns personal information about B—that he has gambling and drinking problems. Can we accuse A of invading B's privacy? 'No,' answers Parent. While it is understandable to feel reticence at calling A's action an *invasion* of privacy, *pace* Parent, it seems unreasonable not to concede that, at the very least, B has *lost* some privacy.

Imagine that A gives the information about B to a reporter who then re-publishes it in a popular contemporary magazine. Isn't B's privacy diminished in this case—doesn't he lose privacy? Again, Parent answers 'No,' on the grounds that it cannot be an invasion of privacy because publicly available information can be found by anyone, 'without resort to snooping or prying' (1983a, 271). But surely one can be made to lose privacy even in the absence of snooping or prying. A flatmate can accidentally walk in on one while one is naked in the shower and see one's private tattoos, for example. In this case, even if it is questionable whether there has been an invasion of privacy—that is, even if it is debatable whether there can be involuntary or blameless *invasions* of privacy—surely privacy has been *lost*. Losses and invasions of privacy are not the same thing.

Parent's account is too narrow and it paints privacy with too thick a brush. It cannot accurately reflect nuances. For Parent, in summary, there is absolutely no difference in loss of privacy between having intimate information about oneself

published in a blog that no one visits and no one will ever read, and having it published in the front cover of *The Guardian,* where millions might read it.

Moreover, contrary to what Parent believes, having undocumented personal knowledge about oneself possessed by others is not necessary for loss of privacy; one can lose privacy even if all that is known about oneself was already in the public domain. Take the following example, inspired in the real world. Imagine that there is highly sensitive information about an individual called Milan in some police files in Czech Republic. The files suggest that Milan was a police informer during the Communist regime in 1950. After, say, 50 years, the files become declassified and enter the public domain. For eight years, no one reads those files. Then one day, in 2008, someone stumbles upon the files and publishes their content in a prominent Czech newsmagazine, which in turn causes the information to become worldwide front-page news.[2]

When did Milan lose privacy? In 2000, when the files were declassified but nobody had read them (it is plausible to think that there was no one alive who was even aware of their existence), or in 2008, when the files were read and published? It seems more accurate to say the latter. Milan's privacy might have been put at *risk* when the files were declassified, but *loss* of privacy occurred when someone

[2] The protagonist of this unfortunate episode was the Czech novelist Milan Kundera. While there is much controversy surrounding these events, I am assuming here for the purposes of the argument that the accusation made against him is true, although in fact I tend to think it is false. The question of whether wrongdoing should be private is a fascinating issue, but a normative one, and beyond the scope of this chapter. Also, I was not able to verify the exact year in which the files were declassified.

read the files, followed by a more grave loss of privacy when the information was published in a Czech newsmagazine, followed by an even more significant loss of privacy when the news was published worldwide. What this example shows is that the degree of accessibility to information and the number of persons who access that information matter for privacy.

The importance of degree of accessibility is at the heart of the European 'right to be forgotten,' which forced Google to allow the deletion of links that can stigmatise individuals for what they did in the past but is no longer relevant for the public. In 2009, the Spaniard Mario Costeja González realised that when he googled himself, the most prominent result was information about a home-foreclosure notice from 1998. Since then, Costeja had paid his debts. The case was brought to the European Court of Justice. The court ruled that Google should remove personal data from search results on a person's name 'when outdated, inaccurate, inadequate, irrelevant, or devoid of purpose, and when there is no public interest' (Powles and Chaparro 2015). The links that Google deletes from its search engine are still alive, but by not appearing in a Google search they are much less accessible.

Another example of how privacy can be damaged even when the information in question is already in the public domain is the case of aggregation. Internet users leave a trace of public bits of information that include a myriad of details about themselves. When scattered, each bit of information may not be meaningful enough to be considered a loss of privacy. Aggregated, however, they can give an

extremely detailed portrait of a person, including her ideas, buying habits, employment status, medical issues, legal problems, personal interests, sexual preferences, geographic location, and more.

It seems, then, that having undocumented personal information known by others is not necessary for losing privacy because one can lose it through documented personal knowledge. One might wonder why someone like Parent might want to hold such an implausible account of privacy, why someone would want to answer 'no' to the question of whether people can lose privacy with aggregation of public information or the publication in a widely read newspaper of documented personal information. Parent replies as follows:

> An affirmative answer blurs the distinction between the public and the private. What belongs to the public domain cannot without glaring paradox be called private; consequently it should not be incorporated within our concept of privacy. (Parent 1983a, 271)

In what follows I will argue that the private/public distinction cannot justify what privacy is or what ought to be kept private. There is no paradox in saying that some events taking place within the public domain can be called private because there are different senses of the word 'private.'

II. Privacy and the public/private divide

Because of the close etymological link between the noun *privacy* and the adjective *private*, it is only natural to assume that privacy—roughly, that which is kept

hidden by individuals from most other people—will always be linked to the 'private' in the dichotomy between the private and the public. Whether this is true depends on how we divide and conceptualise the public and the private. As I shall argue, if we separate these two spheres as people usually do—on the basis of physical spaces, funding, or roles—it is untrue that what is and should be private always belongs to the private sphere.

The private as a physical space

One commonsensical way to divide the private and the public is to suppose that the private sphere is a spatial zone intimately related to individuals and families that is off-limits, or of no concern, to the government or to other citizens. That it is of no concern to others is meant normatively, not descriptively. A neighbour might be very interested in what you look like without your clothes on behind the shut curtains of your bedroom, but he should not be (or, at the very least, he should not seek visual access to your bedroom). A religious conservative government may want to know what sexual practices citizens consensually engage in while in their bedroom, but again, from a liberal point of view (which I am assuming here, for the sake of illustration), it should not have access to that information.[3] The paradigmatic locus of the private in this sense is the home.

By contrast, the public sphere is a shared zone that is the legitimate concern of both the government and the people, the citizenry in general. The public sphere is

[3] I am assuming here a right to privacy which I will argue for in Chapter Four.

an area that is subject to the authority of the government. Paradigmatic public places are streets, public squares, parks, and government buildings.

One might think, then, that in light of this division, people can expect privacy in the guarded comfort of their houses (as long as they do not engage in criminal activities). In the streets, however, privacy is not to be expected or respected. Thus, there are those who think that photographers do not have the right to take shots of people who are in their homes. As soon as people step out of their houses, however, it is fair game to capture them with a lens. Photographer Nick Turpin, for example, believes that 'what happens in a public place should be a matter of public record.' Turpin acknowledges that taking photographs may have bad consequences: 'I could be photographing a couple kissing while they shouldn't be kissing. But if they are doing it in a public space, it's a risk that they're running…' (cited by Laurent 2013).

Turpin's example seems to provide strong support for his position because kissing on the street appears to be a particularly voluntary act. But imagine an individual getting hit by a car on the street. When the paramedics arrive, they cut through his clothes in order to give him medical attention. In this case it is much less clear that witnesses are entitled to take photographs of him naked and injured—much less publish them. That he did not intend to be unclothed in public and the sensitive nature of the situation are reasons for people to act with discretion. Cases such as these are enough to show that Turpin's view is wrong, that people can

have a claim that others respect their privacy when being in a public space—at least in certain circumstances.

Even in cases where it seems that people have chosen to carry out a sensitive act such as kissing in public, it is worth noting that often our 'choices' are much more constrained than they seem. A couple may not have the financial means to pay for a hotel room or live in a place where they can enjoy privacy. A political dissenter may not be able to meet with a fellow activist in her home because she suspects her house is under surveillance. Sometimes, the only breaths of privacy we can get are precisely in public spaces, away from people who can recognise us.

In this vein, Patricia Meyer Spacks notes how women often find privacy only outside of the home: 'the housewife wants privacy specifically to get away from her family for a time' (2003, 1). The home is thought to provide a temporary relief from the demands of the larger society. Granted—at home, one can stay in one's nightgown, put one's feet on the table, and relax. Sometimes, however, domestic chores and family obligations can get in the way of one's individual privacy. Having to engage with family can be an obstacle to spending time by oneself writing in one's diary, for example, or having the chance to speak privately with a friend, away from the eyes and ears of one's spouse and children. With the passage of time, intimacy can become a burden, as mentioned in Chapter One, and caring can become oppressive (Moore 1984, 277, 42). If the house is too small, noises from others' activities can be annoying. An escape from others is important for wellbeing. Men, women, children, and teenagers all get privacy

from their families by leaving their house for some hours every day and entering the public sphere. School, work, and solitary walks (even if they take place in crowded streets) can all serve as potential privacy resources.

Above and beyond these counterexamples, there is an equally important limit to be noted with respect to the public and the private as normative spatial zones. One problem with separating the private and the public on the basis of that which is the legitimate concern of individuals and families, versus what is of concern to the government and the citizenry in general, is that it is a normatively laden demarcation. Such a characterisation could not explain why something should be protected by privacy; it would be question-begging. In other words, it is not enough to say '*x* should be protected by privacy because *x* is part of the private domain.' One would have to then justify why *x* is off-limits or not a legitimate concern of the government or the public.

The private as that which is not publicly funded

Another common way of drawing a line between the private and the public is by appealing to the source of funding. If an institution is funded with money earned by businesses or individuals acting in a non-governmental capacity, it is usually thought to belong to the private sector. According to this way of dividing the pie, individual citizens fall into the private domain on account of their wages, estates, and shares being personal assets. If an institution is funded with money coming from taxpayers, it is usually listed as belonging to the public sector.

One may think that privacy has nothing to do with the source of funding, but it is plausible to argue that citizens are entitled to know about the lives of public officials because they are paying for their salaries by way of their taxes. Thus, one might think that, just as employers are entitled to some degree of knowledge about their employees (at least while the latter are on duty), so citizens are entitled to keep an eye on their public officials.

Financially separating the private and the public is problematic. There are many private businesses (e.g., pharmaceutical companies) that receive public funding or tax breaks and yet are still considered private companies. Private funds (not in the form of taxes) also go into politics, which is traditionally considered part of the public sphere. Most political campaign spending in the United States is privately funded, for example. For Howard Radest, the distinction between private and public (according to standard criteria that include sources of financial support) has eroded and is no longer plausible (1979, 288-289). Although he focuses on the United States, the same point can be made about many, if not all, countries.

Even if we could sharply distinguish between private and public funding, marking the divide by focusing on the origin of money will not help us in deciding what ought to be kept private. Because both individual citizens and private businesses are obligated to pay taxes, it can be argued that all personal finances should be accessible to any other citizen—as is the case in Norway—or at the very least to the eyes of the government. The demand for a certain amount of financial

transparency from all citizens shows that not all that is financially private should be guarded by privacy.

Conversely, publicly funded institutions ought to enjoy privacy at least sometimes. One might worry that lack of transparency might encourage corruption and wrongdoing. However, it is a plausible hypothesis that, at least in some cases, publicly funded institutions might need a degree of opacity. Consider a publicly funded institution that protects human rights—it might need privacy to protect victims and its sources of information. In governmental contexts, when congressional committees are closed, people feel more at ease to speak their minds as well as to compromise with opposing factions without feeling they lose face (Westin 1970, 45). I will look more deeply at the topic of transparency and privacy in institutions in Chapter Six.

The implication of these counterexamples is that what was true of the spatial delimitation between the private and the public is also true of a financial delimitation: privacy cannot be descriptively specified or normatively decided by reference to what is private in a financial sense.

The private as a role we play

A further popular way to distinguish the private from the public is by making reference to the different roles people can play in society. A private role is, roughly, a role in which an individual is entitled to act (and, arguably, in some

cases should act) in her own interest, in the interest of her family, or in the interest of a small group of people to which she belongs. For the sake of argument, I will assume that the paradigm of a public role is that of a public official. I realise, nonetheless, that there are public roles, such as that of celebrities, which may function differently. The role of a public official, in contrast to a private role, is one in which the person should act for the sake of the citizenry as a whole—with *their* interests in mind.

It is a matter of controversy exactly what motivation or method public officials should act on. Some might think that the role of a public official is to act on the basis of the aggregation of individual preferences (as derived from certain institutional mechanisms), independently of the content of those preferences. Others might think that the public official should rather follow her own conscience and do what she sincerely believes will have the better outcomes for society. We can argue about that, but what is clear in any case is that the public official should never act, in her capacity as a public official, for her *own* benefit. As a public servant, she is supposed to serve others. How to best serve the public is a matter outside the scope of this chapter.

People who play a public role also play private roles, but they are expected to separate those two when they are on and off duty. When a Senator is acting as a Senator, she ought to seek to further the interests of the citizens in her society, and the partiality she might feel towards her friends and family should be put in parenthesis when voting about legislation. When that same person goes on

holidays, however, she is allowed to take her public official hat off and exercise her private role as wife or mother or sister, which requires that she should make decisions for the benefit of her family.

Conduct that would be pejoratively called 'favouritism' or 'nepotism' in the context of a public role is considered acceptable behaviour for paradigmatic private roles (Jones 1984, 608). The shopkeeper who is the sole owner of his shop and hires his nephew instead of looking for the most qualified candidate, may not be doing what is best for his business, but he is not doing something morally condemnable, and his action is certainly not comparable to the wrongness involved when a public official takes advantage of his position to hire a nephew for public office. Sometimes private and public roles conflict, and good systems are ones that prevent or minimise such conflicts.

The categorisation proposed is admittedly imperfect. Some roles will not fit neatly into one box or the other. Even if the characterisation is a rough one, however, it is enough to show that even in clear-cut cases of people performing public roles, it is not evident that privacy is out of place. Consider the following example. A public official makes a call to another public official to talk about a matter of public importance. The call is made from a government office and using a mobile phone paid by the government. Although it is all happening within the public realm, and the topic can be of interest to the citizenry, it does not follow that journalists are entitled to spy on the conversation and publish its contents in the press. If such a thing were to happen, as it did in Mexico recently with Lorenzo

Córdova (Glum 2015), the privacy of the public officials involved would be infringed. We may want to justify such violations of privacy in the case of whistleblowers who denounce grave offences, but they still count as privacy infringements, even if justifiable ones. It is clear, in any case, that not *everything* that goes on in the life of a public official while she is on the job should be scrutinised by the public.

Conversely, not everything that goes on when a public official is off duty ought to be private. For example: it is said that, in order to belong to the elite Oxford drinking group, the Bullingdon Club, it is required of current new members that they burn a £50 note in front of a beggar as an initiation ritual (McTague 2013). Suppose that, in thirty years time, a member of today's Bullingdon Club becomes Prime Minister. Even if he was not in public office when he was initiated into the Club, information regarding that event is arguably relevant enough that the future Prime Minister may not have a claim to privacy with respect to it. While the information may not speak to his abilities as a Prime Minister, it speaks to his character and his attitudes towards the underprivileged. As Prime Minister, he will be pushing for legislation that can significantly affect the worst off, and British citizens are entitled to know what his attitudes are towards the poor.[4]

A similar argument is often used when exposing the infidelities of public officials. I believe the argument fails in most cases of infidelity. The Bullingdon Club

[4] My thanks to Cécile Fabre for this example. It should be noted that people may change—particularly when decades have passed. But having such a precedent would put the burden of proof on the politician in question to demonstrate he has changed his attitudes towards the poor.

initiation is politically relevant because it directly speaks to the person's attitude towards the worst off, who form an important section of the population. An infidelity, on the other hand, is usually not politically relevant, as long it is not expressive of unacceptable views about women, for example, or there is no rape or criminal accusation involved.[5]

Some people might want to argue that infidelity suggests that the public official in question is dishonest. Empirical studies suggest, however, that honesty is not a robust trait: a person likely to be honest in one situation can likely be dishonest in a different context; someone can be honest with one's constituents and dishonest with one's family (Hartshorne and May 1928).

Some readers might disagree with the way I characterised private and public roles because my paradigmatic public roles are highly politicised. It is common to think that people are to be separated into private citizens and public figures (or celebrities), depending on the kind of exposure they have to the public in everyday life, the interest they inspire. On this view, it is often believed that public figures have to accept as part of their job description that they are entitled to less privacy than regular people. This terrain is highly controversial, with many grey areas. Among other relevant factors to take into account, we should distinguish between

[5] There may be other cases where it is not clear that an affair is politically irrelevant: for example, when a politician who has voiced homophobic views is found out to have homosexual relationships. Or when the infidelity is committed by a politician who has strongly defended family values and is in a position of power that could contribute to legislate family relationships. Or when politicians get involved with influential people in politics or business who could give rise to conflicts of interest. In most cases, however, the sexual lives of politicians are not politically relevant and should not be exposed.

people who become famous involuntarily (e.g., the victims of a terrible and highly publicised crime), and people who have striven for fame and depend on fame for their success (e.g., actors). It seems reasonable to think, however, that even the most public of figures, the most famous celebrities who want to be famous, are entitled to *some* privacy. We all need some moments of relaxation in order to maintain an acceptable level of wellbeing, and even celebrities ought to enjoy some space and time when they can take a break from their public personae and engage in personal activities that everyone likes to enjoy away from the gaze of others.

Celebrities' claim to some degree privacy is enough to show that this way of parsing the private/public distinction is not enough to decide what should be private. This is not to say that someone being a public figure is not an important consideration when it comes to deliberating about more thorny cases involving claims to privacy on the street. Someone being a public figure, though, is not enough of a justification to invade that person's privacy in any and all circumstances.

III. Conclusion

If we want a definition of privacy that accords or at least resembles our common linguistic usage, social norms, and expectations, and a defensible normative stance regarding what ought to be kept private, it seems that the distinction between the public and the private cannot help us determine either what privacy is or what

should be private. How we characterise the private can have important implications for privacy, of course, but our social norms regarding privacy are much more nuanced and contextual than the broad brush strokes of the distinction between the public and the private.

Realising the divide between privacy and different senses of the private, Hannah Arendt argues that '[t]he distinction between the private and the public realms, seen from the viewpoint of privacy rather than of the body politic, equals the distinction between things that should be shown and things that should be hidden' (1998, 72). While this way of separating the public and the private is accurate when it comes to privacy, it is also tautological. It cannot serve as a way to justify what should be private, as one needs to give an explanation of why something should be hidden or shown. (Not that Arendt meant this characterisation to be justificatory of privacy.)

I do not mean to argue that the dichotomy between private and public spheres is without any value. If one's project is to work out whether and when the government may interfere with individuals' freedom, or if one wants to find out which checks and balances are appropriate for someone who administrates public funds (as opposed to private ones), then different ways of distinguishing the public and the private may be useful. What I have argued here is merely that the distinction between private and public as it is usually construed is not as useful as it may seem at first glance to determine what privacy is and what ought to be protected by privacy. There are many different senses of what the private is, and

no paradox is involved in saying that one can have a private conversation in a public square.

CHAPTER THREE

Eight Accounts of Privacy—And Their Shortcomings

In the previous chapter I went through a very popular but misguided way of thinking about privacy. In this chapter I turn to academic debates and go through some of the definitions and accounts of privacy that are most influential in the literature in law and philosophy—those that turn on the concepts of the right to be left alone, control, possession of information, and limited access, as well as the attempt to understand privacy as a 'family resemblance' concept and the turn to reductionism. For each account I will point out its strengths and weaknesses.

This chapter serves as preliminary work for the next chapter, where I present my own definition of privacy and the right to privacy.

I. (1) Privacy as being let alone

In 1890, Samuel Warren and Louis Brandeis published their famous article "The Right to Privacy," which constitutes the first theoretical attempt to lay out the nature and importance of privacy. Warren and Brandeis were responding to the development of mass media, and thought of privacy as an issue related to journalists publishing private information and photographs:

> Instantaneous photographs and newspaper enterprise have invaded the sacred precincts of private and domestic life; and numerous devices threaten to make good the prediction that "what is whispered in the closet shall be proclaimed from the house-tops" (Warren and Brandeis 1890, 195).

Warren and Brandeis recognised privacy as entrenched in culture and evolving over time in response to social and technological developments. They believed that '[p]olitical, social, and economic changes entail the recognition of new rights,' and that 'the common law' should develop to 'meet the demands of society' (1890, 193).

Warren and Brandeis thought of privacy as intimately related to being let alone. Their view is particularly understandable given that their target was the yellow press. Journalists who take pictures invasively and print gossip are a nuisance—it is not for nothing that they are called *paparazzi*, an Italian dialect word that describes the exasperating noise of a buzzing mosquito.

Although the matter is somewhat controversial, it appears that Warren, a lawyer from Boston, was especially sensitive to the 'overstepping' of the press on account of having married the daughter of a senator. From the time Warren got engaged to Mabel Bayard in 1882 until "The Right to Privacy" was published in 1890, there were almost sixty articles about the Bayard-Warrens in the newspapers (Gajda 2008, 44). The content of the articles included comments about the 'accentuated' hips of the new Mrs. Warren, information about her spending a lofty amount of money on a painting, and detailed descriptions of the deaths and

funerals of two family members (Gajda 2008, 45, 41). It is no wonder that Warren and Brandeis would later argue that the law should 'protect those persons with whose affairs the community has no legitimate concern, from being dragged into an undesirable and undesired publicity' (1890, 214).

Warren and Brandeis, then, argued that the right to privacy is an instance of the right 'to be let alone':

> the protection afforded to thoughts, sentiments, and emotions, expressed through the medium of writing or the arts, so far as it consists in preventing publication, is merely an instance of the enforcement of the more general right of the individual to be let alone. It is like the right not to be assaulted or beaten, the right not to be imprisoned, the right not to be maliciously prosecuted, the right not to be defamed. (…) The principle which protects personal writings and all other personal production, not against theft and physical appropriation, but against publication in any form, is in reality not the principle of property, but that of an inviolate personality.[1] (Warren and Brandeis 1890, 205)

Many philosophers (e.g., Davis 2009, Parent 1983b, Thomson 1975) have interpreted the cited passage to mean that an 'instance' refers to an example or single occurrence of the right to be let alone. According to this interpretation, someone has privacy just when he is let alone and loses privacy when he is not let

[1] To speak strictly, Warren and Brandeis do not define privacy. At most, they attempted to define the *right* to privacy, though the bulk of their article is about defending the importance of privacy. It can be argued, however, that one can infer a definition of privacy from their definition of the right to privacy. If, according to Warren and Brandeis, the right to privacy is an instance of the right to be let alone, then we have reason to think they thought of privacy as an instance of being let alone. Most privacy scholars have taken this interpretation to be true. It is important to acknowledge nonetheless that there is a significant degree of speculation on this inference. As my own account will show, privacy and the right to privacy may be two very different things. A number of scholars have endorsed this definition of privacy as being let alone. See, for example, (Bloustein 1978, 123-186, Freund 1971, 182-198, Posner 1981, 314-316).

alone. Under this common interpretation, this account captures both too much and too little.

As Judith Jarvis Thomson points out, if we hit Jones with a brick, we have not let him alone, yet it does not seem that we have violated his privacy (we seem to have violated some other right) (1975, 295). If someone can keep his privacy intact while not being let alone, the implication is that letting someone alone is not necessary for that person to retain his privacy.[2]

Conversely, as Steven Davis points out, someone who voluntarily discloses personal information by sharing it out loud in the village square is being let alone and yet is losing privacy with respect to the information he conveys to the villagers who are walking by (2009, 451). This example shows that letting someone alone is not sufficient for him to retain his privacy. Another example is the use of methods of surveillance of which the victim is not aware. In one sense, it seems that if we do not touch the person or go near her but only watch her every online activity on the Internet, we let her alone, in the sense that we let her do as she pleases without

[2] Warren and Brandeis's account is the first but not the last theory to come up in law that seems at best partial and at worst far from our common understanding of privacy. I suspect part of why accounts in law can be so distant from common sense understandings is that, more often than not, law theorists are attempting to offer accounts that will protect people's privacy within existing legislation that was not originally designed to protect privacy. Warren and Brandeis drew on threads of past jurisprudence and tried to construct a legal concept of personality using property doctrine, tort law, copyright law, and damage principles. They trace protection of privacy by the law of contract, trespass, defamation, and breach of confidence to argue that courts could protect privacy without the need to legislate a new cause of action in tort.

interference, yet it seems on the face of it that we are at the same time invading her privacy.[3]

In my view, a more accurate and fair interpretation of Warren and Brandeis's phrasing is to understand privacy as a *subcategory* within the *general* right to be let alone (Gavison 1980, 437, footnote 48). If this interpretation is correct, then they might have meant to defend the view according to which all people who are let alone have privacy, but not all people who are not let alone do not have privacy. This is because one can fail to be let alone by virtue of someone violating some other subcategory of the general right to be let alone (different from privacy). If this is the case, then the objection that letting someone alone is not sufficient to secure privacy is still pertinent, but not the objection about necessity. In other words, hitting Jones with a brick while respecting his privacy is not really a counterexample because we are not letting him alone in virtue of violating some other kind of right under the broader heading of the right to be let alone (probably the right not to be assaulted). In contrast, the fact that someone may lose his privacy while being let alone (as in the village square case) is still a counterexample, because under this interpretation, all people who are being let alone also have privacy.

[3] Thomson offers a similar example: 'The police might say, "We grant we used a special X-ray device on Smith, so as to be able to watch him through the walls of his house; we grant we trained an amplifying device on him so as to be able to hear everything he said; but we let him strictly alone: we didn't touch him, we didn't even go near him—our devices operate at a distance"' (1975, 295). I chose Internet surveillance as an example because of its contemporary relevance and because X-ray devices physically harm individuals through radiation, even if they cannot feel it.

More importantly, this account of privacy is unsatisfactory because it suffers from underdescription. Warren and Brandeis do not tell us what is special about privacy—what makes privacy privacy and not some other kind of right under the general heading of being let alone. Being let alone is the genus, and having privacy is a species of being let alone, according to this interpretation. It might be true that when one's privacy is being invaded, one is not being let alone, but more would have to be said in order to differentiate *privacy* invasions from other kinds of invasions. Privacy cannot consist *just* in being let alone.

Moreover, as presented by Warren and Brandeis, the concept of being let alone is so vague that it can be easily manipulated to fit people's interest. The authors do not define what they mean by being let alone. The person who screams out intimate information in the village square might try to argue that she is not being let alone because there are too many people on the square who are staring at her. Similarly, intelligence agencies like the NSA could argue that they are respecting people's privacy because they are letting them alone, in the sense that their surveillance is unobtrusive.

In summary, the definition of privacy as being let alone falls into two pitfalls: it is too restrictive (it is not sufficient for privacy), and it is underdescribed.

II. Control-based definitions

(2) Control of information

Charles Fried has defined privacy in terms of control over information:

> Privacy is not simply an absence of information about us in the minds of others; rather it is the control we have over information about ourselves. (…) The person who enjoys privacy is able to grant or deny access to others. (Fried 1970, 140)

Randall Bezanson agrees:

> The idea of privacy in the 1990s expresses more clearly the individual's interest in some measure of control over self through control over information. (Bezanson 1992, 1140)

Rather than focus on publication in newspapers, as Warren and Brandeis did, Bezanson thinks 'regulation should focus on the collection of information and access to it, with regulation at the point of publication serving only a secondary protective function' (1992, 1136).

There are problems with defining privacy on the basis of control, but I will mention those below, when I discuss other control-based definitions. For now, I will only mention one important problem that arises from restricting privacy to information.

Consider the case of a Peeping Tom who catches a glimpse of a naked victim. Intuitively, it seems that privacy is not only about information, but also about

physical access. The sheer gaze of someone can feel like an invasion of privacy, as was mentioned in Chapter One, even when that uncomfortableness cannot be explained through the information that the watcher gains. I will consider an objection to this view later on in this chapter.

This account also covers too little. A more promising account of privacy as control, then, might be one that goes beyond information and includes other kinds of access.[4]

(3) Control over being sensed by others

For Richard Parker, privacy amounts to having control over who can sense us:

> *[P]rivacy is control over when and by whom the various parts of us can be sensed by others.* By "sensed," is meant simply seen, heard, touched, smelled, or tasted. By "parts of us," is meant the parts of our bodies, our voices, and the products of our bodies. "Parts of us" also includes objects very closely associated with us. By "closely associated" is meant primarily what is spatially associated. The objects which are "parts of us" are objects we usually keep with us or locked up in a place accessible only to us. In our culture, these objects might be the contents of our purse, pocket, or safe deposit box, or the pages of our diaries.[5] (Parker 1974, 281)

Parker's focus on the physical has the implication of making personal information and privacy come apart. Though the previous discussion showed how privacy

[4] Privacy has also been characterised in terms of the control a person has on information about him or herself by Alan Westin (1970), Arthur Miller (1971), and Elizabeth Beardsley (1971), among others.

[5] Other philosophers who contend that privacy is about control over a certain realm of life include: Elizabeth Beardsley (1971), Robert Gerstein (1978), James Rachels (1975), Jeffrey Reiman (1976), and Richard Wasserstrom (1978).

cannot be explained adequately only through personal information, doing away with sensitive information and focusing *only* on sensorial access seems just as unreasonable.

The following examples show the counterintuitive implications of Parker's definition. Parker argues that '[i]f we tell someone that we are homosexual, we lose control over private information, but we do not necessarily lose privacy' because '[w]e choose to let him hear us, and retain the power to stop speaking to him' (1974, 282, 294). It seems to me that it makes more sense to say that when we share sensitive information with close ones, we relinquish privacy in exchange for intimacy, understanding, acceptance, and meaningful connections. Similarly, Parker contends that 'if X read Y's diary, Y's loss of privacy would be the same regardless of the contents of the diary. The offense against Y's privacy is that the diary was seen, meaning in this case, read' (282). These implications are unpalatable and they point to two sources of problems.

The first source is the focus on physical access to the detriment of personal information. For Parker, it is the exact same loss of privacy if someone reads an empty diary that we just bought and have not had time to make use of—a diary devoid of personal information—as if they read our life-long diary where our darkest secrets are kept, because 'in either case, [we] lost the same degree of control over who saw [our] diary,' (1974, 282) and 'the information revealed when we are sensed does not affect the degree of the loss in privacy' (283). However, if most of us were given the choice to relinquish an empty diary or a full one to

someone else, I am quite certain most people would choose the first option, probably without even feeling *any* loss of privacy had occurred (perhaps a loss of autonomy, but not of privacy).

Parker believes that people respond strongly to their information being violated because the collection of information about people 'devalues privacy' by making the individual lose control of the flow information about himself, by making privacy less secure (threatening privacy), and by making the individual 'constantly uncertain' of whether he still has privacy (285). Parker's explanation seems like a convoluted way of accounting for what is undesirable about having someone access one's personal information. It seems inaccurate to describe a situation where someone acquires a great deal of sensitive information about ourselves through other means than sensing us (e.g., through other people's gossip) as one in which we still have privacy, but it is a 'devalued' kind of privacy. It seems to me that privacy is something generally valuable to us, and a devalued privacy is no privacy at all.

The second source of problems for Parker is the focus on control. The problem is seen most clearly when he discusses the possibility of someone, call him Hicks,[6] recording a conversation we are having at a party. Parker says that: '[w]hether the recording is ever replayed has no effect on the degree of loss of privacy, for the loss consists not in being listened to, but in losing control over when and by whom one

[6] 'Hicks' was the codename of the famous British radio producer and spy, Guy Burgess. For illustrative purposes, throughout this chapter I will use this codename as a proxy for someone who spies, records, eavesdrops, etc.

is listened to' (283). In other words, for Parker there is no difference in loss of privacy between the following possibilities: Hicks having a recorder at a party but not using it (Hicks having the power to record me and me not having control over the recorder, and in consequence not having control over who might listen to me), Hicks recording our conversation but never replaying it and keeping it in a safe, Hicks recording our conversation and replaying it to himself (assuming here that I had the conversation with Hicks), Hicks recording a conversation I had with someone else and replaying it to his friend, and Hicks recording a conversation I had with someone else and broadcasting it on the BBC radio.

A satisfactory account of privacy should be sensitive to the degrees of loss of privacy this gradation of examples shows. A control-based account is unlikely to be able to do so, because as soon as someone has a recorder that is not in our power, we have lost control of who might hear what we say. Julie Inness, however, thinks that control-based accounts of privacy can respond to this objection.

(4) Control over intimate decisions

Julie Inness argues that

> privacy is the state of possessing control over *decisions* concerning matters that draw their meaning from the agent's love, liking, and care. In other words, claims to privacy are claims to possess autonomy with respect to love, liking, and care. (Inness 1992, 140) [Emphasis added]

The kind of decisions Inness has in mind are decisions about who we allow to have intimate access to ourselves (she includes informational access as a subset of

access (63), as well as sensorial access), and other decisions about intimacy such as 'child rearing and education, family relationships, procreation, marriage, contraception, and abortion' (64). Inness equates love, liking, and care to intimacy. She believes that

> To call *Y* [an intimate decision about *x*] "intimate" is to claim that it involves a choice on the agent's part about how to embody her love, care, or liking. *Y* involves such a choice because *x* [range over instances of access, instances of information dissemination, and the agent's activities] draws its meaning and value from the agent's care, love, or liking. (Inness 1992, 91)

It is unclear what Inness means by 'care.' At no point in her book does she define the term, but her theory seems to suggest that she has in mind a relational concept of caring for *someone* (attending to them, taking care of them, being fond of them), rather than a more abstract caring, as in 'I care about what happens to my country.' When discussing the value of privacy, for example, she says that 'privacy is valuable because it acknowledges our respect for persons as autonomous beings with the capacity to love, care, and like—in other words, persons with the potential to freely develop close relationships' (95). And elsewhere she says that 'the attitudes of love, liking, and care are directed toward the development of affiliation with others' (87).

A major problem with a conception of privacy focused on what is shared with our close ones (those whom we love, like, and care (for or about)) is that it is incapable of explaining privacy in regard to personal information that someone might not want to share with *anyone*—not even her loved ones. In fact, it is plausible to argue that there are some kinds of personal information which people might be

particularly interested in hiding from those to whom they are closest. That Inness's theory is incapable of dealing with these cases is exemplified in her discussion of a legal trial: *Melvin v. Reid.* In this case, a movie had been made depicting the past of the plaintiff as a prostitute and using her maiden name. The plaintiff won the suit. Inness analyses this case in the following way:

> *Melvin* involved privacy considerations because the defendant disseminated *intimate*, personal details about the plaintiff's former life as a prostitute, thus allowing others to gain *intimate informational access* to the plaintiff's life. Such details about an agent's past life are intimate because they are commonly imbued with emotional significance as far as their sharing is concerned; typically, we share secrets about our past with those for whom we feel love, care, or liking. (Inness 1992, 128)

Inness does not consider the possibility that the biggest privacy damage in this case might have been the disclosure of sensitive information to Melvin's own family. Melvin might have been adamant in hiding her past particularly from those she loved, cared, and liked the most for fear of feeling rejected or misunderstood, or to protect her children from the potentially traumatising truth that their mother had to become a prostitute to earn a living. Often our darkest secrets are more easily shared with strangers than with our families. Because we care more about the people we love, we are particularly vulnerable to them, and they to us (if love is mutual). To protect ourselves and our nearest and dearest, in some cases we hide painful truths from them. It is not uncommon for veterans, for example, not to want to speak to those whom they love, care, and like about their experiences in war—what they saw, what they did, and what was done to them.

Likewise, Inness's theory cannot explain why most people feel it is an invasion of privacy when a stranger goes through their dustbins. Rubbish is not the kind of thing we typically share with those whom we love and like.

A further questionable element in Inness's proposal is her inclusion of private decisions such as childrearing and abortion as privacy issues. She justifies it by appealing once again to the intimacy of such decisions:

> We commonly distinguish between intimate and nonintimate decisions about our actions, characterizing intimate decisions as "private" or "personal"—unfit subjects for the state's regulatory power. Consider the difference between being informed that the social welfare mandates that we must engage in sexual activity with specified individuals and being informed that the social welfare mandates that we must pay taxes. (Inness 1992, 64)

Liberty of action is curtailed in both instances, but it is only morally acceptable in the latter case. Inness's choice of examples is a good one, and we might be persuaded that personal decisions should be included in the realm of privacy protections if we stay with the examples she provides. As soon as we leave those extremes, however, it becomes much less obvious what should be considered personal or intimate decisions. We might think that what characterises personal decisions are sensitive themes such as sex and the body. Indeed, Daniel Solove, in his justification for the inclusion of what he calls 'decisional interference' in his taxonomy of privacy, offers the following observation:

> Decisional interference and exposure have been judicially recognized to affect the same aspects of the self—health, the body, sex, and so on. The decisional interference cases track traditional areas that are widely considered to be private, such as the home, family, and body. (…) Decisional interference [also]

> bears a similarity to the harm of intrusion as both involve invasions into realms where we believe people should be free from the incursions of others. (Solove 2006, 559)

In a similar vein, Judith DeCew defends the inclusion of decisional interference in privacy due to the '*nature* of the decision[s],' which 'involve issues related to one's body, family relations, life style, or child rearing' (1986, 165, 159-160).

We deceive ourselves if we think the government does not and should not get involved in matters related to our bodies or families (both realms typically considered as intimate spheres). Obligatory schooling is one way in which the government heavily influences childrearing and gets involved in family matters; child abuse is prohibited within (and outside of) the family; maybe more controversially, polygamy is banned in most countries in the world. Any government is bound to have regulations on the body and the family, and it does not seem to be a matter of privacy. On this point I side with critics such as William Parent (1983a, 273) and Louis Henkin (1974), who have pointed out that decisional interference cases primarily concern issues of autonomy and liberty, not privacy. As Richard Posner writes:

> we already have perfectly good words—liberty, autonomy, freedom—to describe the interest in being allowed to do what one wants (or chooses) without interference. We should not define privacy to mean the same thing and thereby obscure its other meanings. (Posner 1981, 274-275) [7, 8]

[7] It seems to me that, once again, we face a case of distortion of a concept produced by the legal point of view—by lawyers' laudable desire to protect people. Some lawyers believe that people should be free to decide, for example, whether to have an abortion or not. From this point of departure they will use any and all legal tools available to them to defend their case before judges and juries. In some cases, the recourse to privacy has succeeded in convincing judges and

There are some decisions about one's life or one's family's life that should only be made by individuals themselves, without any interference from the government. There are other decisions about one's life or one's family's life that may be justifiably influenced by the government. The line between these two spheres is constantly disputed and negotiated, and it is geographically, culturally, and historically contextual. The separation in question, however, is not delineated by the home, the body, the intimate, or the private, as my examples show. Furthermore, it should be noted that the defence of 'decisional privacy' seems like a particularly normatively laden one, as it seems to imply which kinds of decisions *should* be up to the individual, and perhaps not ideal to define the concept of privacy as such.

It is possible that the confusion of considering intimate decisions a matter of privacy arises because these kinds of decisions involve personal information that is related to privacy. For example, if the government regulates contraception, it will likely need to have personal information about individuals that it should not be allowed to have for privacy reasons. But the privacy concerns here lie in the information, not the decisions themselves; the decisions themselves are a matter of reproductive rights and freedom. Another possible source of confusion is Warren and Brandeis's legacy of thinking of privacy as an instance of being let alone.

---

juries that women should be free to decide whether to have an abortion, but that does not mean it is a conceptually valid move.

[8] To return to the etymology of the word *privacy*, there is no evidence that anything like so-called 'decisional privacy' was part of the early meaning of the term, which further supports criticisms against this notion.

More generally, apart from the *content* of privacy (whether it is about information, sensorial access, or intimate decisions), my main criticism of all three control-based definitions of privacy (those of Fried and Bezanson, Parker, and Inness) is that control does not seem to be either a necessary or a sufficient condition for privacy.[9]

One can have no control over one's information and sensory accessibility and yet enjoy privacy. Scientists at Harvard University have recently developed a drone the size of an insect that has been dubbed RoboBee (Poole 2013). Future generations of RoboBees can be expected to carry minuscule video cameras for surveillance; call this prototype RoboSpyBee. Suppose, for the sake of argument, that there is a RoboSpyBee already developed and sitting in a scientist's lab at Harvard. It seems to me that, even though RoboSpyBee is outside of my control, my privacy has not been lost just because the Harvard scientist has acquired the power to invade my privacy.[10] As long as RoboSpyBee is not following me around, my privacy remains intact. A further example is given by William Parent: a comatose patient may have no control over his privacy, but if others protect it for him (or if others ignore him), he cannot be said to be devoid of privacy (1983b,

[9] Another eminent privacy thinker worth mentioning who is an advocate for a control-based definition is Alan Westin, who believes '[p]rivacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others' (1970, 7).

[10] Judith Jarvis Thomson also makes this point: 'It is the actual looking that violates [someone's privacy], not the acquisition of power to look' (1975, 305, footnote 1).

343-344). If one can lack control and yet enjoy privacy, control is not necessary for privacy.

Conversely, one can have unlimited control over one's information and sensory accessibility and yet have no privacy. Suppose I tell something sensitive about myself to someone. In that moment, I lose some privacy with respect to him, and some control (as he may divulge that information). But suppose further that I incarcerate that individual and allow him no contact with the outside world. In this case, I seem to retain control over my personal information, but it is still true that I have lost some privacy to my prisoner, which implies that control is not sufficient for privacy.

On the sufficiency claim, it can also be argued that if someone decides to divulge personal information about himself to the public (through a blog, the press, TV, etc.), it seems that he loses privacy as he exercises control over his information. Politicians know that if they want to become public figures, they must surrender a part of their privacy to the mass media. They regularly do it voluntarily. It seems, then, that one can lose privacy precisely by exercising control over one's private information and personal access.

The most plausible counterargument I have found in the literature available to those who believe control is the essence of privacy is the following. Inness argues that the second kind of objection, that one can lose privacy while exercising control of one's personal information,

> fails because of its overly simplistic notion of control. (…) Exercising control is an ongoing process; as such, it consists of not only the voluntary initiation of a situation, but also the ability to regulate the situation as it develops (which includes the ability to either continue or halt it) and a reasonable expectation of continued control. Furthermore, an agent must be able to exercise this regulative ability with respect to her desired end, rather than an arbitrary or imposed end. (Inness 1992, 48-49)

Inness further clarifies what she means by control: 'We have control in situation X when there exists a reasonable probability that we could regulate the outcome of the situation without recourse to emergency maneuvers' (51).

If one decides to disclose intimate details about oneself to the public, explains Inness, one *can* be said to lose privacy under a control-based conception of privacy because, even if one had control to initiate this action, the moment one gives information to strangers, one loses the ability to regulate its future dissemination, because strangers 'do not possess duties to conceal information that has been revealed to them' (1992, 55). In contrast, if one discloses information to friends, though one forgoes 'physical control' of the information, one still retains 'equally effective control due to social norms and ties of friendship' (55, footnote 18).

Inness's notion of control, however, seems too demanding to be realistic. Rarely can we regulate the initiation, development, and future of anything in life, privacy included. Control is an illusion. Inness wants to have us believe that we retain control of our information, and hence privacy, when we confide in friends. The pervasiveness of gossip, however, defies this idea. As soon as we share information with others, we lose control because we cannot recall it, we have lost the option of

keeping it to ourselves, we cannot erase it from others' minds.[11] Furthermore, as soon as we speak about intimate issues, our information might not even be in our friend's control: she might involuntarily talk about it in her sleep; unbeknownst to both, someone might be overhearing us, or an NSA analyst may be listening to and recording our conversation through one of our mobile phones. But perhaps we are lucky. Maybe no one is overhearing or recording our conversation, and maybe our trusted and trustworthy friend does not talk in her sleep and takes our secrets to the grave. If this is the case, then we have only lost privacy with respect to our friend, and we have suffered no privacy invasion because it was a voluntary act. Even if we lack control of that information once we share it, our privacy with respect to people other than our friend can remain intact, which further shows that privacy cannot be defined on the basis of control.

III. (5) Privacy as keeping personal information safe

Inspired by William Parent's account of privacy, which was explored in detail in the previous chapter, Steven Davis (2009) has made a crucial observation: that many attempts to define privacy mistake having a *right* to privacy for having *privacy* itself. Borrowing heavily from Parent's account, he offers the following definition of privacy:

---

[11] One could object that this point is in direct contradiction with my criticism that when we share sensitive information with friends, we are losing privacy while exercising control. The apparent contradiction points to the equivocation in the term 'control'—it is not clear whether it is meant to refer to the exercise of autonomy (in which case one does exercise control when sharing information with friends), or the ability to restrain the flow of information or access (in which case one does seem to relinquish control when sharing information with a friend). Inness seems to refer to the latter.

> In society *T*, *S*, where *S* can be an individual, institution, or a group, possess privacy with respect to some information, *p*, and some individual, *U*, if and only if:
> a) *p* is personal information about *S*.
> b) *U* is not in an informational state with respect to *p* nor is the information, *p*, readily available to *U*.
> In society *T*, *p* is personal information about *S* iff (…) most people in *T* would not want anyone, other than him/herself, to be in an informational state with respect to *q* where *q* is information about them which is similar to *p*, or *S* is a very sensitive person who does not want anyone, other than him/herself, to be in an informational state with respect to *p*. In both cases, an allowance must be made for information that most people or *S* make available or would make available to a limited number of other people or to a certain subset of people. (Davis 2009, 455)

Davis inherits from Parent the advantage of specifying that privacy is about *personal* information. He further strengthens his view by explaining how privacy comes in degrees: one can have privacy with regard to a certain bit of information and a certain individual. The more individuals know about a bit of personal information about ourselves, the less privacy we have with respect to that information. A further advantage of his view is that it can accommodate different cultural conceptions of what is to be considered private.

It could be argued that it is unwise to make an account of privacy so culturally-relative because it prevents us from criticising exhibitionist trends. For example, younger generations seem to share much more information about themselves on Facebook than what older generations deem appropriate. If that amount of sharing becomes the norm, then much information previously thought of as personal will cease to be so, so that we will not be able to say of people who share intimate information on Facebook that they will be losing privacy. We can,

however, criticise the culture itself. There is still much room for normative considerations. We can give reasons why certain kinds of information should be considered personal, and why they should not be widely shared. We can point out that the domain of privacy of younger generations is narrower than that of other generations, and we can give reasons for why that is desirable or undesirable.

Davis's account is not without problems, however. He argues that informational states include knowing and believing that *p* (as long as the belief is warranted), as well as non-propositional information involving sensory modalities (e.g., how someone looks). He does not say, however, how one is to judge whether a belief is warranted. And even if we could agree on criteria for a belief to be considered warranted, it is unclear how one can lose privacy due to a false, yet warranted, belief.[12]

Take the following example. Suppose I have an office mate, Lenina, who seems to have gained weight lately. I also notice that she is unusually giddy, and that every morning she spends an inordinate amount of time in the bathroom. As a result of these oddities, I believe Lenina is happily pregnant and suffering from morning sickness, but out of discretion, I tell no one. Let us agree that the belief is

[12] Davis is not clear about whether he wants to include false beliefs in 'informational states.' He criticises Parent's definition for making knowledge a necessary condition for the loss of privacy and argues that in some cases 'there can be a loss of privacy for someone when others come to have personal information about him that they believe rather than know' (2009, 454). But he does not discuss the case of false beliefs; he only gives an example of a warranted belief that happens to be true. If he is only willing to entertain justified true beliefs as causes for privacy losses, then it could be argued that the former simply collapses into knowledge, which would imply that Davis ends up sharing Parent's view on this point.

warranted, for the sake of argument. And let us assume it is true that most people in our society would not want others to know or believe they were pregnant if they themselves had chosen not to reveal that information. In fact, Lenina is not pregnant but instead has a new boyfriend who is a chef (hence the gaining of weight), and whom she calls every morning from the bathroom in order to have privacy. Davis's account would have us think that Lenina lost privacy on account of my false belief in her pregnancy, even if I did not publicise that belief. It is unclear to me whether Lenina lost any privacy at all. If she did, however, it is not because of my believing something false about her, but because of me noticing something strange about her and perhaps paying closer attention to her than what is normal, which would typically lead me to learn more from her than if I had not paid attention.

A further problem is that Davis's account equates someone possessing information with someone having ready availability to that information. It seems, however, that one has less privacy if Hicks reads one's emails or diary than if he merely has readily available access to them but never reads them, as I indicated when I discussed Parker's view. The second situation might also be bad for privacy, but there is an important nuance in degree that Davis's account does not capture.

Furthermore, Davis's rendering of 'ready availability' is rather implausible. He gives an unpalatable example to illustrate what he means by something being readily available:

> It is easy for employees at my university's computing service to read my e-mail. This, however, does not make my e-mail readily available to them, as I understand the notion of ready availability. There are thousands of people who use my university's e-mail system and out of these, the employees have not decided that they are going to read my mail, although it would not be difficult for them to do so. Since they have not thought about reading my e-mail, it is not readily available to them. Had they thought about reading my e-mail, I would suffer a loss of privacy. (Davis 2009, 457)

It seems far-fetched to make losses of privacy depend on others' thoughts about whether they will access one's personal information or not. If we forget our diary in a friend's house, but she never reads it, it seems that we have not lost any privacy, regardless of whether our friend had a fleeting desire to read it at some point.

Finally, Davis's account does not seem to give a satisfactory account of physical sensorial privacy. Consider the case of a reoffending Peeping Tom who already knows what Victim looks like unclothed, but who continues to peep. It would seem that every time Tom peeps, Victim loses privacy, even if Tom does not acquire any new knowledge or beliefs. Davis thinks that his definition does account for this case. The information, argues Davis, *is* different every time Tom peeps because it is indexed with respect to the exact time at which he gains the non-propositional information. Different informational states corresponding to two different times of peeping have different truth conditions. This observation seems beside the point, however.

Imagine the Peeping Tom looks through Victim's window for a whole hour. He learns nothing new because Victim is going through a daily routine that Tom has seen many times before. Presumably the truth conditions for every second Tom peeps through the window are different (i.e., is he looking at Victim at t_1? At t_2? At t_3?), but this fact is absolutely irrelevant for Victim. Suppose Victim knows that Tom has seen him many times before. Victim is not worried about Tom's *informational* states, much less about their truth conditions; if he were to explain the wrongness of what Tom does, he would likely never make reference to those factors. Victim is made uncomfortable not only by the propositional and non-propositional information that Tom may acquire about him, but also by Tom's sheer gaze. There is an almost physical discomfort in having someone look at one when one does not want to be looked at that does not seem to be captured by describing it in terms of the *information* others may acquire about one.

IV. (6) Privacy as limited access

Ruth Gavison agrees that the realm of privacy is broader than that delimited by information. She believes privacy is better described as a concern over our accessibility to others:[13]

> In its most suggestive sense, privacy is a limitation of others' access to an individual. As a methodological starting point, I suggest that an individual enjoys *perfect* privacy when he is completely inaccessible to others. This may be broken into three independent components: in perfect privacy no one has any

[13] Other philosophers who contend that privacy is about restricted access (of various kinds) include: Anita Allen (1988), Roland Garrett (1974), Hyman Gross (1971), and William Parent (1983b).

> information about *X,* no one pays any attention to *X*, and no one has physical access to *X.* (Gavison 1980, 428)

On this account there are three components of such limited access: secrecy, anonymity, and solitude.[14] We lose privacy when others obtain information about us, when others pay attention to us, or when others gain physical access to us.

One may wonder what unifies these three concerns. For Gavison, we claim privacy in these three different ways for similar reasons (i.e., privacy has a series of functions in our lives): 'the promotion of liberty, autonomy, selfhood, and human relations, and furthering the existence of a free society' (1980, 423). More succinctly, privacy allows individuals to do what they otherwise would not do for fear of others' reactions in different spheres of life (friendship, politics, academia, etc.).

The first category proposed by Gavison, *secrecy*, seems slightly inapposite. While the limiting of personal information is certainly a crucial aspect of privacy, as the limitations of Parent and Davis's accounts suggest, it is too broad to say that privacy is about information in general. It seems that not just *any* kind of

[14] Gavison thinks of physical access as the kind of 'physical proximity' that permits *Y* to observe, touch, and hear *X* through the normal use of his senses (1980, 433). The ability to watch and listen through technological devices at a distance is thus not included in this category. It is not clear to me why Gavison chooses to exclude sensing the person at a distance. Later on in her paper she says that her concept of privacy covers invasions of privacy such as photographing individuals (436). She must think, then, photographing individuals as part of either informational privacy or attentional privacy.

information can be related to privacy. Barring strange situations, the colour of my doorbell does not seem like a private piece of information.[15]

The second problem with this category is that 'private' and 'secret' are not always coextensive. Not everything that is private information is a secret. Consider genetic information. While we may not want to publish the results of a genetic test online for fear of what others might do with it (e.g., perhaps a future employer who is racist would not want to hire one if one has the 'wrong' genetic background), it would be strange to consider it a secret. Suppose one had never even looked at the results of the genetic test. It would still be private information about oneself, but it could hardly be considered a secret if the person owning the secret (and whom the secret concerns) does not even know the content of the secret. Who could be said to be keeping the secret from whom? Conversely, not everything that is secret is private information. As Inness points out, the organisation of a surprise party may be a secret to the birthday person, but it is not usually considered private information (1992, 60).

Gavison's second category, *anonymity*, is likewise problematic. She argues that a person can lose privacy by being paid attention to (by being followed, stared at,

[15] We can imagine a circumstance in which the colour of one's doorbell is significant for privacy. Imagine that someone who knows that Victim has HIV publishes an article on the newspaper saying that there is a person who lives in Victim's street that has HIV, and that that person lives in the house with the red doorbell. The bit of information about the red doorbell takes on a whole new significance in this context, and suddenly counts as *personal* information because it has become the kind of information that people would typically not want to divulge widely.

listened to, or observed in any other way). She argues that attention is different from information, even though these two ways of accessing someone are related:

> Attention is a primary way of acquiring information, and sometimes is essential to such acquisition, but attention alone will cause a loss of privacy even if no new information becomes known. This becomes clear when we consider the effect of calling, "Here is the President," should he attempt to walk the streets incognito. No further information is given, but none is necessary. The President loses whatever privacy his temporary anonymity could give him. He loses it because attention has focused on him. (Gavison 1980, 432)

Later on in her article, Gavison affirms that 'the aspect of anonymity that relates to attention and privacy is that of being "lost in a crowd"' (1980, 434, footnote 40). It seems to me that what Gavison was after is 'inconspicuousness'—the quality of going by unnoticed. Somebody can be the centre of attention in a crowd (say, because of the way he is dressed) and still be anonymous, in the sense of not being identified by others as any particular person with a name, an address, and so on. More importantly, the stress of being singled out in a crowd can be explained away through the other two categories contemplated by Gavison: informational and physical access. When someone calls 'Here is the President!' in the middle of a crowd, people who are there receive relevant information (the President is here at this time; that person there is the President),[16] and they attain physical access to him (they probably turn to look at him). It seems, then, that the notion of losing privacy through being the focus of attention is unnecessary,

[16] While this kind of information ('here is this person') is not typically considered private or personal information, depending on the context, it might be. If one is supposed to be elsewhere, it may suddenly become sensitive information. Also, one may argue that where one is at a time is personal information when one is famous because it is the kind of thing nobody would want to share (if they were famous).

because such cases can be explained by people gaining informational and physical access to one.

Gavison's third category, *solitude*, seems to correctly identify an aspect of privacy, even if the label she chose for it is again somewhat misguided. It is often true that seeking privacy is coextensive with seeking solitude. Privacy, however, is not always individual. Couples typically seek privacy now and then, and we would not call it solitude. It follows that when we want privacy, we do not always want to be alone, but we do want to be able to choose our company. When we are in situations where we feel vulnerable (e.g., in a hospital bed), we typically want to be able to choose who has physical access to us. In any case, the limitation of physical access is a crucial aspect of privacy, of which the Peeping Tom is a paradigm. The case of the Peeping Tom, however, shows that sensorial access at a distance should be included within this category. Whether the Peeping Tom is looking through our window or our webcam, what is of most concern is the visual access he gains to us.

Independently of the aspects of privacy she proposed, Gavison's account has been criticised for her focus on the limitation of access. Parent claims that limited access is neither necessary nor sufficient for privacy (1983b, 345-346). One could imagine *X* not having privacy with respect to some information despite access to *X* being limited. Consider Hicks eavesdropping on a conversation through a closed door. Access can be significantly limited: let us imagine Hicks cannot make out every word, and when cars go by outside the building, he ceases to hear; but now

and then he can hear phrases that give him an overall sense of the conversation. Another example, given by Parent, is that of a policeman who wiretaps a person despite access being limited by the requirement of getting permission from a judge before listening in (1983b, 346). These examples show that limited access does not seem sufficient for privacy. Conversely, one could imagine someone having unlimited access to *X* and yet respecting *X*'s privacy. Consider a person at a changing room in the gym who could look at *X* naked but chooses to look away or turn around. Or imagine someone who stumbles into someone else's diary but does not read it. These examples purport to show that limited access is not necessary for privacy.

Gavison, however, could respond by arguing that the policeman has effectively transcended the previously limited access by getting a warrant and has achieved full access once he wiretaps someone's phone. Similarly, if Hicks can make out the sense of the conversation he is eavesdropping on, Gavison could argue that it follows that he has (enough) access to that conversation to invade the interlocutors' privacy. Likewise, Gavison could say the person who looks away in a gym is effectively limiting her access.

What is missing from Gavison's account, however, is further specification of the appropriate limits to access that are to count as privacy. It is not enough to say that access is limited. Gavison has in mind a gradation: from perfect privacy (e.g., being in a cave in a deserted island) to no privacy at all (e.g., being naked and plugged into a mind-reading machine that broadcasts thoughts). But both

situations are impossible in life. In the real world, access is almost always limited (i.e., we still cannot read each other's minds), but there are some limits that are much more relevant for privacy than others (e.g., underwear is usually more important for privacy than outerwear). Gavison cannot account for the difference in relevance that limits have for privacy. If we follow her definition, as soon as some passer-by catches our eye in the street, we lose privacy; but this seems too trivial. Talk of limited access in general, thus, is not specific enough to be informative, accurate, or useful to think about privacy. We must look for *relevant* limits to access.

V. (7) Reductionism

A natural reaction to the difficulties of coming up with an adequate definition of privacy is to think that perhaps there is no unified definition of the term that can make sense of our intuitions. Maybe all the connotations ascribed to privacy are hopelessly heterogeneous. Some philosophers have thus argued that trying to define privacy is misguided, and that the concept of privacy can be explained away in terms of its composing elements without any loss.

One way to go about this is to focus on the kind of right we should be protecting. In an article entitled "The Right to Privacy," Judith Jarvis Thomson defends what she dubs the 'simplifying hypothesis':

> [T]he right to privacy is itself a cluster of rights, and (…) it is not a distinct cluster of rights but itself intersects with the cluster of rights which the right over the person consists in and also with

> the cluster of rights which owning property consists in.[17] (Thomson 1975, 306)

In other words, according to Thomson, we may want to call this heterogeneous cluster 'rights of privacy,' but these rights lack a common foundation. Each is in fact an instance of some other, more fundamental, right. If we want to protect privacy, she argues, what we have to do is protect those other rights.

Thomson starts by presenting two paradigmatic cases of violations of privacy. The first scenario, The Pornographic Picture, involves someone spying on someone else's home:

> Consider a man [(call him Bernard)] who owns a pornographic picture. He wants that nobody but him shall ever see that picture—perhaps because he feels that nobody shall know that he owns it, perhaps because he feels that someone else's seeing it would drain it of power to please. So he keeps it locked in his wall-safe, and takes it out to look at only at night or after pulling down the shades and closing the curtains. [Hicks has] heard about his picture, and [he wants] to see it, so [he trains his] X-ray device on the wall-safe and look[s] in. (Thomson 1975, 298-299)

To own a picture, Thomson argues, implies that one has a variety of positive rights with respect to it: one can sell it, tear it, modify it, stare at it, etc. It also implies that one has negative rights with respect to it: a right that others will not take it away and, more controversially, a right that others not look at it, among other rights. She thinks that, even if we have the right that others not look at our belongings, we might not always want to claim that right. Sometimes we invite

[17] Similarly, H.J. McCloskey argues that 'any right to privacy will be a derivative one from other rights and other goods' (1980, 35).

others to look, and sometimes we are just not careful enough (within reasonable limits) to protect our belongings from others' glances. In those cases, if someone looks at our belonging, no right is violated because we have 'waived' our right (302-304, 311).

The second scenario, The Quiet Fight, is a case of eavesdropping. Suppose a couple are

> having a quiet fight, behind closed windows, and cannot be heard by the normal person who passes by; and suppose that (…)[Hicks] trains an amplifier on [their] house, by means of which he can hear what [they] say; and suppose that he does this in order to hear what [they] say... (Thomson 1975, 296)

Thomson believes that, while The Pornographic Picture story involves a violation of property rights, The Quiet Fight involves a violation of rights over the person. She argues that bodily 'ownership,' or self-ownership, gives rise to a right to limit others' access in the same way property ownership generates a right of inaccessibility. We have a right that others not touch us, see us, hear us, etc. (Thomson also mentions the right that we not have our words transcribed, and that we not be modelled in bas-relief as further examples.) We can, however, waive these rights—voluntarily or through some kind of negligence—by letting others touch, see, hear us, and so on. If a couple is screaming while they argue, and a passer-by hears them, no right of theirs is being violated because, by yelling, the couple waives their right not to be heard. Their rights would have been violated had they taken 'conventional and easily available steps' to prevent

listening (e.g., closing the windows, lowering their voices), and Hicks had used his amplifying device to listen in to them (1975, 306).

For Thomson, then, every privacy right is overlapped by other, more fundamental, rights and can thus be fully explained in terms of the right from which it is derived (property rights and rights of the person). According to her, 'it is possible to explain in the case of each right in the cluster how come we have it without ever once mentioning the right to privacy' (1975, 313).

The strength of Thomson's argument lies in its parsimony—if successful, it does away with an allegedly useless concept—as well as in the valuable contribution made by calling attention to the close relationship between ownership and privacy. It is indeed true that access to others' possessions or person often constitutes a violation of privacy. Contrary to what Thomson argues, however, having a justified property or self-ownership claim is neither necessary nor sufficient for having a privacy claim. To illustrate this point, we can return to the case of The Pornographic Picture. Imagine that the picture in question is not really Bernard's, but Mustapha's. Suppose Mustapha lent it to Bernard for a period of time. Furthermore, suppose that Bernard is not actually in his house but is rather in Helmholtz's (a friend) vacation cabin, where he has locked the picture in a safe also owned by Helmholtz. Suppose, then, that Hicks, our spy, turns on his X-ray machine to Helmholtz's cabin to see Mustapha's pornographic picture (which is temporarily enjoyed by Bernard). It is not at all obvious whose privacy Hicks is invading. It seems reasonable to argue that Helmholtz's privacy is being

invaded, as Hicks is getting a glimpse of his home, a very personal possession. However, it is equally reasonable to say that Bernard's privacy is being invaded (particularly if Hicks sees Bernard enjoying the picture), even though Bernard owns neither the picture nor the cabin. Furthermore, it is also reasonable to argue that Mustapha's privacy is not being invaded, particularly if Hicks has no knowledge of Mustapha's existence, even if he is the rightful owner of the picture.[18]

Consider another example: it seems reasonable to think that Fanny has a right to privacy that other people (say, in the train) not look at the passport she carries in her purse (suppose that she does not want people to see it because she hates the way she looks in her passport picture), even if her passport is in fact not owned by her, but by the government that issued it.

As a third example, consider dustbins. Items discarded in the rubbish are usually considered abandoned property—meaning that the owner has relinquished her claim to the contents of the bin. Yet Hicks could learn a lot from someone just from taking her rubbish and analysing it. He could plausibly get DNA samples, he might find intimate items such as pregnancy tests, he would know what his victim buys and eats, he might even find papers that convey his victim's thoughts and

[18] In a similar vein, Julie Inness offers a good counterexample to Thomson's view: 'imagine that I have written a number of love letters to another person. By sending these letters to the person, I relinquish possession of them, yet although the letters are no longer mine, my privacy is still violated if my lover copies the letters and distributes them to others without my consent' (1992, 33).

feelings (discarded drafts of essays or a memoir), as well as financial information about his victim (e.g., by reconstructing shredded bank statements).

These three examples show that ownership is not necessary for privacy; that is, one can lose privacy even if one does not own the items that one's privacy protects. Ownership is also not necessary for having a *claim* or a *right* to privacy, as it seems we have a right that persons not snoop around things that would make us lose privacy even when we do not own those things.

Ownership does not seem to be sufficient either for someone to have a right to privacy. It is not obviously true that owning a bicycle automatically gives one the right that others not look at it. If we had a right that others not look at anything we own, then the counterintuitive implication would be that we are constantly waiving dozens of rights every time we step out of our homes. The proliferation of questionable rights seems more unpalatable than the difficulties and possible shortcomings of coming up with satisfactory definitions of privacy and the right to privacy. Importantly, Thomson's theory has no way of explaining the difference between a diary and a classroom notebook or a novel, between a picture of someone naked and a magazine ad. If privacy rights are nothing above and beyond ownership rights (I include here self-ownership), then everything owned is *equally* private, regardless of the nature of the object, and every time someone sees us, hears us, or smells us, we are waiving batches of rights.

Admittedly, the notion of 'personal information' *is* prominent in the paper, though Thomson neither defines it nor points out its importance. In a discussion on torture and extortion by threat, she contends that we would violate a right of privacy (apart from violating other rights, such as the right not to be hurt) only if the information we were after was 'personal' (1975, 308). When discussing whether there is a violation of some right if someone publishes in a newspaper that Bernard owns a pornographic picture, Thomson argues that the violation could be explained by

> our having a right that others shall not cause us distress, and then add that what is violated here is the right to not be caused distress by the publication of *personal* information, which is one of the rights which the right to privacy consists in, and one of the rights which the right to not be caused distress consists in.[19] (Thomson 1975, 309) [Emphasis added]

The 'simplifying hypothesis' may not ever mention the word 'privacy' to explain why we have a right, but if it is to be minimally plausible it does seem that it has to surreptitiously rely on concepts such as 'personal information' that intuitively are at the heart of what we mean by privacy. Finally, even if a right to privacy is a subtype of ownership rights, Thomson has failed to give an account of what marks out the distinct kind of ownership right the right to privacy is.

[19] Thomson does not fully support the view she rehearses in this quote because she believes that 'even if there is a right to not be caused distress by the publication of personal information, it is mostly, if not always, overridden by (…) the public's right to a press which prints any and all information, personal or impersonal, which it deems newsworthy' (1975, 310).

VI. (8) Family resemblance

There are three strategies to try to give a satisfactory account of privacy. The first is to attempt to define it through necessary and sufficient conditions (accounts 1-6). The second is to appeal to reductionism, arguing that privacy or the right to privacy can be better understood as dissolving in terms of other, more basic, concepts (account 7). The third proposes to understand privacy not as a unified concept, but as a cluster of different meanings, all glued together through resemblance. Daniel Solove defends the view that

> privacy is better understood by drawing from Ludwig Wittgenstein's notion of "family resemblances." As Wittgenstein suggests, certain concepts might not have a single common characteristic; rather they draw from a common pool of similar elements. (Solove 2002, 1091)

Solove believes that it is not possible to come up with a single definition that will share a common denominator and include all types of privacy. According to him, '[p]rivacy is a concept in disarray. Nobody can articulate what it means' (Solove 2006, 477). The best we can do, he argues, is to come up with a taxonomy of privacy: a list of privacy-related items that justify their inclusion in the cluster by dint of their resemblance to other items on the list.

For Solove, privacy amounts to 'protection from a cluster of related activities that impinge upon people in related ways' (2006, 484). His taxonomy organises such problematic activities into four groups:

1. Information collection, which includes surveillance and interrogation.

2. Information processing, which encompasses aggregation, identification, information insecurity, secondary uses of collected information, and failure to communicate to data subjects the information that others have about them.
3. Information dissemination, which includes breaches of confidentiality and disclosures, blackmail, appropriation (identity theft), and the dissemination of false information about individuals.
4. Invasion, which is comprised by intrusion ('invasive acts that disturb one's tranquility or solitude'), and decisional interference (Solove 2006, 491).

Solove points out various resemblance relations. Here I review some of them. For Solove, interrogation is similar to intrusion in that it is invasive. 'Like disclosure, interrogation often involves the divulging of concealed information' (2006, 501). He does not explain what he means by 'divulging.' If he means to say that disclosure is like interrogation in that interrogations often lead to involuntary disclosures on the part of the person being interrogated, then it seems to me a mistake to say they are similar: that one may lead to the other does not imply resemblance. If he means to say that disclosure is like interrogation in that the content of information in both cases is often communicated to a wider audience, then he seems to be confusing the attaining of information with the dissemination of information.

Solove goes on to argue that surveillance resembles interrogation in that both involve gathering information without the consent of the subject of information.

Both identification and disclosure involve the revelation of true information. 'In some ways, identification resembles interrogation, as identification often involves the questioning of individuals to compel them to identify themselves' (2006, 514). Once again, Solove seems to confound two practices often going hand in hand with two practices being similar to each other. It is true identification is often achieved through interrogation, but this does not imply resemblance. Identification, furthermore, can be achieved through other means (e.g., aggregation of information). It is like saying that making documentaries is similar to doing interviews, as documentaries often include interviews. But there are many documentaries that do not include interviews. Solove uses the same kind of reasoning when it comes to intrusion and disclosure, and once again when it comes to decisional interference and blackmail: 'Intrusion is related to disclosure, as disclosure is often made possible by intrusive information gathering activities' such as surveillance and interrogation (2006, 553). 'Decisional interference also bears an indirect resemblance to blackmail, in that laws restricting consensual private sexual behavior often give rise to blackmail' (2006, 561). Finally, according to Solove, decisional interference resembles intrusion in that 'both involve invasions into realms where we believe people should be free from the incursions of others' (2006, 559).

As Hans Sluga (2006) has successfully argued, the notion of family resemblance borrows from two very different ways of understanding the relationship between words, which correspond to different vocabularies: that of kinship, and that of similarity. From the perspective of kinship, families are families in virtue of causal

connections of descent, and things are part of the same umbrella term (e.g., 'games,' 'privacy') in virtue of a linguistic history. If families depend on causal connections of descent, then similarities in family members 'do not constitute that membership'—'they only provide evidence for it' (Sluga 2006, 15-16).

Given that Solove points out supposed similarities among the cases he includes in his taxonomy, it is unlikely that he is appealing to kinship relationships, and it is wise of him not to. While kinship terms cannot be defined once and for all because they are always open to new members—we never know who will part of the family or what will count as privacy in the future—they can be defined by 'limiting ourselves to a particular moment, or to how these terms have been applied in the past' (Sluga 2006, 18). Solove wants to say that privacy cannot be defined at all, so it makes sense for him to understand family resemblance as relations of similarity.

What makes similarity terms indefinable is that 'the range of relevant similarities is not fully determined' (Sluga 2006, 18). Solove himself briefly entertains the objection that such open-endedness can imply that similarity concepts are boundless, endless. If that were true, similarity concepts could include anything at all, as one can always find *some* resemblance between any two items—even if they are among the most disparate of things (e.g., dogs and mortgages resemble each other in that they can be found in planet Earth). As a result, the concept of privacy would be empty—*anything* could be tied to it. Solove responds to this objection as follows:

> Although Wittgenstein suggests that not all conceptions are "closed by a frontier," this does not mean that conceptions are endless. Rather, it means that not all conceptions have fixed and sharp boundaries separating them from other conceptions. Boundaries can be fuzzy or can be in a state of constant flux. (Solove 2002, 1098)

Solove misses the target. His response does not address the fundamental issue of how to make sure concepts are not endless, even if they have fuzzy edges or change. Another way to phrase the challenge is to ask what is needed for a resemblance to count as a *family resemblance*.

Michael Simon answers this question by suggesting that 'two or more things may be said to fall within the range of a family-resemblance term if and only if they can be connected by relevant similarities to a single paradigm' (1969, 412). On this view, one cannot say that *a*, *b* and *c* are *x*s just because *a* resembles *b* and *c* resembles *b*. Rather, one would need to find a paradigm case *p* and argue that *a*, *b*, and *c* are *x*s in virtue of their resemblance to *p*. One implication of this proposal is that it is possible that *a*, *b*, and *c* may not resemble each other in any (non-trivial) way and still be *x*s.

Even this more restricted version of family resemblances encompasses too much. Even if we could find a way of discarding all trivial resemblances, such as being found in planet Earth, there are other resemblances that are non-trivial and yet seem inappropriate for the justification of membership. Take surveillance as a paradigm case of a privacy issue. One could argue, for instance, that boxing blows, blood transfusions, and torture are all privacy-related issues because they

resemble surveillance in that they are all invasive. But that is clearly false, under any common use of the word *privacy*.

If all words were thought of in terms of family resemblances, it would be hard or impossible to distinguish one term from another, as all terms would encompass an inordinate amount of things.

The problem of unboundedness in defining privacy spills onto ethics. How we define privacy will have an impact on what we think should be protected by privacy. If there is no way to select what kinds of things belong to the concept of privacy, we will have no clue as to what kinds of things ought to be protected by privacy.

VII. In search of a better definition

Given the difficulty of proposing a satisfactory definition of privacy, some authors decide not to define it at all. In *On Privacy*, Anabelle Lever warns that her book 'will not advocate or defend a particular definition of privacy, on the grounds that this task—if necessary or useful—may be easiest to accomplish once we have a better sense of the philosophical terrain involved' (2012, 7).

There is some wisdom to this strategy. People talk about privacy all the time without first defining it; we often have an intuitive notion of the term we can work with. Pretty soon, however, a lack of a definition is an obstacle to meaningful

reflection and debate. When Lever advocates for the protection of privacy, for example, it is hard to know what she means exactly.

It seems to me that it is better to have a definition of privacy that one knows is not perfect (because it may be controversial, or it may leave out a few cases that should be included, or incorporate others that should not be included) than no definition at all. Having a definition allows others to assess our claims with precision. It also allows for progress to be made. If one compares old accounts of privacy, such as that of Warren and Brandeis, with more recent ones, one gets a sense that definitions of privacy have gone through a process of perfection and refinement through time. It is plausible to argue that we have a better understanding of privacy today than we did in 1890, and hopefully we will have yet a better understanding in the future.

In the following chapter, I offer my own definition of privacy, which is intended to contribute to the improvement of our understanding of the concept. It builds upon some of the virtues of the definitions I have reviewed here and tries to avoid the pitfalls pointed out in this chapter.

CHAPTER FOUR

Privacy, the Right to Privacy, Perceptions of Privacy, and Privacy-related Obligations: A Map of the Moral Territory

As we saw in Chapter One, people seek privacy to escape the burdens and risks of social interaction. We take refuge in our bedrooms, our offices, and away from the gaze of others to be safe and unperturbed—to cry, to recover, to explore, to relax. Sometimes we give up privacy voluntarily and in exchange for other goods we value, as when we share our secrets with our friends in return for closeness and understanding. Sometimes privacy is taken away from us without our consent, as when we are spied on.

In this chapter, I distinguish between privacy and the right to privacy. I offer a definition of privacy and contend that it is a good we enjoy thinly (i.e., in actual circumstances). The right to privacy, in contrast, is a right to a robustly demanding good (i.e., it has counterfactual or modal demands). When we enjoy the right to privacy, we enjoy an assurance that our privacy will be respected, not only here and now, but also in relevant possible worlds (e.g., independently of our income, or our political or religious views). I argue that one may have privacy and yet have one's right to privacy violated. Conversely, one may have one's right to privacy respected while losing privacy.

My definition of privacy is meant to be a descriptive account. It is concerned with actual circumstances and objects of privacy (e.g., personal information) and it is supposed to help us answer questions about whether an individual has lost privacy or not. My account of the right to privacy is a normative one. It is concerned, not only with actual circumstances, but also with counterfactual ones, and not with the objects of privacy, but rather with the ways of getting at objects of privacy (e.g., spying as a way to get personal information about someone). It is supposed to help us answer questions about whether an individual has been wronged with regard to her privacy.

Section I offers some notes on terminology. Section II establishes the account of rights I will be using. Section III proposes a definition of privacy as remaining personally unaccessed. Section IV explores the value of privacy and the interests we have in keeping certain things private. Those interests ground the right to privacy, understood as the claim to robust privacy, which is the focus of Section V. Section VI clarifies the role of social norms for privacy and the right to privacy. In Section VII, I argue that perceptions of privacy should be taken into account in the moral exploration of privacy cases. The chapter ends with Section VIII, where I sketch some moral reasons to protect one's own privacy.

I. A word on terminology

According to my terminology, privacy can never be violated, only the right to privacy can be violated—privacy can only be lost or invaded. A privacy loss can be voluntary or involuntary (both on the part of the subject of privacy and on the part of the person who gains access to another's privacy). A privacy invasion happens when a person is made to lose privacy without her consent; when it is unjustified, it amounts to a violation of the right to privacy. However, the two are not synonymous because, as will become apparent, some unjustified violations of the right to privacy do not entail an invasion of privacy.

Regarding rights that protect robustly demanding goods, I talk about *failing to respect rights* and distinguish that from *rights violations*. Take security as an example. In order for individual *I* to fully respect subject *S*'s right to security, she must be unwilling to harm *S* here and now and in relevant possible circumstances. If *I* is not harming *S* out of indifference or laziness but *would* harm her if she had the chance or if *S* did something that displeased her, she fails to respect *S*'s right to security, even if she has not yet violated *S*'s right to security. Failing to respect a right is not as morally grave as violating it.

I violates *S*'s right to security only if she

> a) secures a position from which to harm her and with the intention of harming her now, in the future, or in some counterfactual situation (e.g.,

plants a bomb in *S*'s house that can be detonated should *S* do something that displeases *I*), or

b) attempts to harm her (e.g., tries to detonate the bomb but fails due to some technical problem), or

c) succeeds in harming her (e.g., successfully detonates the bomb and hurts *S*).

Harming a victim is morally worse than attempting to harm her, as it involves a rights violation *and* a damage (the loss of the good of physical integrity). Attempting to harm a victim is usually morally worse than securing a position to harm her. Attempting to harm someone is typically just the next step from securing a position that will allow one to carry out that attempt. More importantly, an attempt usually subjects the victim to a higher risk of being harmed, and the perpetrator usually believes this to be true (otherwise, presumably, she would not carry out the attempt).

I will now sketch out the account of rights that my arguments will presuppose.

II. More on rights

Raz's Interest Theory

I do not intend to advance a general theory of rights here, as it is beyond the scope of this work. I understand rights roughly along the lines of the Interest Theory

proposed by Joseph Raz (1988, 1984).[1] On this view, someone has a right if she can have rights—if her wellbeing is of ultimate value[2]—and if her interest in having that right is a sufficient reason for holding other people to be under a duty.

Being of 'ultimate value' means being of intrinsic value, above and beyond instrumental value or the value of consequences. Although only those whose wellbeing is of ultimate value can be the bearers of rights, interests do not have to be of ultimate value in order to be the basis for rights. Journalists, for example, have a right to protect their sources even if this interest is valued because of its usefulness to society (Raz 1984, 206-207). The fundamental role of rights in practical thinking, then, is to ground duties in the interests of beings whose wellbeing is intrinsically valuable. In other words, rights restrict others' actions in the interest of right-bearers.

Although rights protect interests, the right-holder need not want to enjoy the right. A right to education, for example, grounds a duty to provide opportunities for education to each individual, even when the individual in question does not desire to be educated. Furthermore, though rights are based on the interests of right-holders, there are cases in which an individual has rights that go against her interests—for example, someone may own a property that brings her great trouble (Raz 1984, 208). This is explained by the fact that rights are bestowed on right-holders on the basis of their *general* characteristics (e.g., being citizens, being

[1] Other proponents of the Interest Theory of rights include Matthew Kramer (e.g., Kramer 2010) and Neil MacCormick (e.g., MacCormick 1977).
[2] Raz also thinks that 'artificial persons' such as corporations can have rights.

creatures with certain needs, etc.), and not on the basis of their personal, individual characteristics.

Rights are agreements to protect people. Raz understands rights to be 'intermediate conclusions' in arguments that go from ultimate values to duties—they stand between interests and duties. Interests partly justify rights, which in turn partly justify duties. We use talk of rights as a sort of shorthand so that, on the one hand, we can save time and not have to refer to ultimate values every time a practical question comes up, and on the other hand, so we can build a common culture that allows the protection of people's interests despite there being confusion and disagreement concerning ultimate values. In other words, we can agree about the rights people have even when we do not agree on the justification for those rights. This agreement turns rights into reasons, even if they are not ultimate reasons—i.e., rights constitute reasons for action even independently of their ultimate justification (Raz 1984, 208-209).

Like Raz (1988, 170-172), I do not endorse strict correlativity between duties and rights—the idea that for every duty there is one corresponding right and vice versa—for the following reasons. First, one right may ground many duties, not just one (Raz 1984, 199, Waldron 1989, 503, 509-512). The right to free speech may not only ground a duty on the part of the government to refrain from supressing speech—it may also ground a duty to protect people who give public speeches from the potential wrath of the mob, another duty to create spaces where people can be heard, etc.

Second, sometimes we may recognise the existence of a right (such as the right to education) without being certain who is bound by duties based on that right (Raz 1984, 211). We can disagree about whether it is society or parents who are the duty-bearers. Similarly, on some occasions it might be clear that we have an obligation to do some action (e.g., because it is effortless on our part and it would make the world a much better place) even if there is no clear right-holder to whom we owe that obligation—even if there is no one who can have a grievance against us for failing to perform that action. In cases where the effort is more burdensome and the benefit less significant, a moral loss might still occur if we fail to perform the good act, but not to the point where we can be accused of wrongdoing.

Finally, rights have a 'dynamic' aspect to them: changes in context and circumstances, such as technological advancements, may lead to there being new duties grounded on the old right, which means duties are as unpredictable as the future (Raz 1984, 200).

Raz's framework of rights supports a view of morality that is not wholly rights-based (Raz 1988, 193-216). In contrast to philosophers like Ronald Dworkin (1997, 185-222), who thinks that political morality is rights-based, or J.L. Mackie (Mackie 1978), who thinks all of morality is rights-based, in Raz's view and my own, morality is not exhausted by rights, as there are moral reasons for action that may not be grounded in rights. For example, one ought to give information to people which it is in their interest to have even when they have no right to it (Raz

1988, 196). In cases where we ought to do something even if there is no right-bearer to whom we owe it to, I will refer to *obligations* (even if Raz does not differentiate between duties and obligations). If failing to carry out the good act does not create enough of a moral loss to be considered wrongdoing, I will write about having moral reasons for action. Rights-based moralities cannot account for the moral value of virtues, the pursuit of excellence, or supererogation (Raz 1988, 196). If rights and their correlative duties exhausted morality, we would not be able to make sense of the moral value of going beyond the call of duty.

Conflicting rights

Raz's account of rights makes it appear as if rights can never conflict. He does acknowledge that there can be conflicting considerations or reasons, however (Raz 1984, 209-210). In those cases, conflicting considerations must be weighed against each other. If *X*'s interests are less weighty than *Y*'s interests, then the would-be right of *X* does not arise, because, given the counteracting conflicting considerations, we cannot hold ourselves to have a duty towards *X*.

In order to determine the existence of a right, argues Raz, there must be premises concerning the interest of the right-holder, the required importance of the interest, and the particular person or class of persons who are obligated to the right-holder. In addition, one needs to add other premises 'establishing that these grounds are not altogether outweighed by conflicting reasons' (Raz 1984, 209). Later on, Raz affirms that '[i]f conflicting considerations show that the basis of the would-be

right is not enough to justify holding anyone subject to any duty, then the right does not exist' (1984, 210). This stance leads Raz to say that '[a] general right is (…) only a *prima facie* ground for the existence of a particular right in circumstances to which it applies' (1984, 211). This view seems unsatisfying for those of us who believe rights are possessed on the basis of the interests of right-holders, and not depending on the external circumstances individuals may find themselves in. If rights are based on the interests of individuals, then those rights should not just disappear when faced with conflicting considerations, as the interests of individuals do not disappear even in the face of conflicts. At most, interests can conflict, but not disappear.

Raz points out that when conflicting considerations only show that *some* actions cannot be required as a duty, other actions may still hold as a matter of duty, in which case the right does exist but is able to ground 'duties only for some of the actions which could promote the interest on which it is based' (1984, 210). Though Raz does not speak of conflicts between rights as such (only of conflicting considerations), he does admit in a footnote in his paper "On the Nature of Rights" that '[c]onflicts of rights are possible if conflicts of duties are,' though he does not further explore these conflicts (1984, 211, footnote 1).

A more detailed and convincing account of conflicts in the context of Raz's Interest Theory, which I endorse here, is offered by Jeremy Waldron in his paper "Rights in Conflict" (1989). According to Waldron, conflicts between rights are inevitable if we understand rights along Raz's lines, because even in the case of a

single individual, interests often conflict. For Waldron, rights conflict when duties are incompatible—when it is not possible for all the duties to be performed. For example, there is a conflict of rights when two people who are drowning have a right to be rescued but there are resources to save only one of them.

Waldron argues that when there is a trade-off of one right against another—when we are forced to comply only with the duties grounded in one right—we should not sacrifice the losing right. It is not that 'a consideration loses the status of a right when it happens to conflict with another,' but rather that we are failing to fulfil our duties with respect to that right (Waldron 1989, 508). Even if there are not enough resources to rescue two drowning persons, both still have a right to be rescued. Waldron's account is more convincing than that of Raz because it is truer to the spirit of rights: if a right is a right, it should not just disappear in the face of a conflict. Waldron argues convincingly that rights can generate successive waves of duties. In the rescuing dilemma, successive duties might impose constraints on the way resources are allocated in a society (with more resources having to go to rescuing services); the training of more professional rescuers might be required; as well as an investigation to diagnose the causes of the shortage of resources; and it may also be required that compensation be given to victims of trade-offs or their families. Thus, even in the case where our primary duty (to rescue both people) is not fulfilled, the rights of people do not disappear, and further duties are created: 'the trading off of one right against another, in a situation of conflict, is never the end of the story' (Waldron 1989, 512).

That rights ground a multiplicity of duties and can create successive waves of duties complicates the act of balancing and trading-off of rights. There is no simple or formulaic way of resolving such conflicts. Suppose we agree that the right not to be tortured is weightier than the right to free speech—we have still to throw into the balance the duty to investigate torture, for example, or to compensate victims or their families, along with similar requirements created by the right to free speech. It may be the case that, given limited resources, marginal duties related to torture are lower in the ranking than urgent action related to free speech and political freedom. So even if the right not to be tortured is more important than most other rights, not every duty associated with it is to be given priority (Waldron 1989, 515).

In some cases, conflicts between rights may be best resolved in a quasi-quantitative way by balancing what is at stake, following the metaphor of weight: once the relative importance of the interests at stake has been established, and consideration has been given to the contribution made by conflicting duties to the interests rights are meant to protect, we try to maximise what we deem of most importance (Waldron 1989, 515).

In some other cases, however, attention to internal connections may be more appropriate to resolve conflicts. Take the case of free speech. If there is a group of people (e.g., Nazis) calling for the suppression of another group of people, such that the right to free speech of both groups can be said to be in conflict, the correct strategy is not to take a quantitative approach and count the number of

people in both groups to try to determine the course of action that will produce the violation of the fewest number of rights, giving the right claim of each individual in both groups the same weight. Rather, if we understand free speech as an interest in participating as equals in a public sphere where all may speak their minds, then we can agree that to count as a true exercise of free speech, assertions must be of a kind that allow for opposing views to be expressed as well. The Nazi speeches only conflict with the free speech of others because they aim to bring an end to the form of life that makes free speech possible. Their speeches are incompatible with the right they are asserting, and should therefore be banned. Thus, in some cases, when we cash out the internal relations between rights claims, what looked like a conflict between rival interests may dissolve (Waldron 1989, 518).

These considerations about rights will be relevant for this chapter as well as the next, where I discuss the conflict between the rights to privacy and security. In what follows I offer a definition of privacy.

III. Privacy as remaining personally unaccessed

Let us take Tom's peeping through Victim's bedroom window as a paradigmatic example of a loss of privacy. This example encapsulates many a lesson about privacy. After some introductory observations inspired by the example of the Peeping Tom, I will offer a proper definition of privacy.

Numbers matter

Other things being equal, one loses more privacy when one is exposed to more people. Victim is less badly off if it is only Tom peeping through her window than if there are more people watching.

Privacy is not an all-or-nothing affair

A lesson drawn from previous accounts of privacy reviewed on Chapter Three is that, even if Victim loses privacy with respect to Tom, she does not lose *all* of her privacy. He might be seeing her as she undresses or he might listen to a personal conversation she is having on the phone, but surely there are still many private matters that are inaccessible to Tom (e.g., the Victim's finances, her love letters, etc.). It is therefore imprecise to say that 'one has privacy' or 'one does not have privacy,' although I will sometimes use this kind of expression for the sake of brevity. If we want to be precise, we need to specify whether we have privacy or not with respect to an object x (where x can be one's naked body, one's personal conversation, one's love letters, etc.) and an individual I (at time t).

Privacy is not only about information

Suppose Victim is reading the newspaper in her bedroom and is not worried about the trivial information Tom can glean by spying on her through the window. Again, as we saw in Chapter Three, she might still be upset about Tom's leer, even when she is not concerned about the information he is acquiring.

The unease experienced when someone looks at one in a certain way or when one does not want to be looked at does not seem to be captured by describing it in

terms of the information others may acquire about one. Indeed, if Victim were to realise that Tom is looking through her window (and supposing there is nothing she can do to make him go away), she could be so disturbed that she might no longer be able to concentrate on her reading.

Culture matters

In most Western cultures, Victim would be understandably upset if a Peeping Tom were to see her breasts as she undresses in her bedroom, and not so upset if Tom saw her knees. In some African tribes, however, breasts are not conceptualised as private parts, and women are considered to be naked when they expose their knees or thighs.

There does not seem to be any reason to think that a conception of female nakedness based on breasts is superior to one that is based on knees and thighs. This example suggests that the kind of things people wish to keep private is at least partly determined by culture, even if there may be some universal categories such as 'nakedness' that are understood differently in different cultures. As I indicated in Chapter One, there seem to be many more commonalities than differences in the kinds of things that people all over the world wish to keep private.

These observations support the following definition of privacy (all technical terms are defined further below):

In society *C*, a subject *S* possesses privacy with respect to *p*, and vis-à-vis some individual *I*, if and only if,

a) *p* is either:

i. personal information about *S*, or

ii. *S*'s *autotopos*

b) and *I*:

has not accessed *p*

For *S* to have complete privacy with regards to *I*, *I* would have no knowledge of *S*'s personal information, and no access to *S*'s autotopos.

In society *C*, *p*(i) is personal information about *S* if either

- it is common for people in *C* not to want anyone, other than themselves (and perhaps a very limited number of other people chosen by them), to know about *p*(i), or
- *S* is sensitive about *p*(i) (even if in his society it is uncommon for people to care about others having access to that kind of information about themselves),[3] or

[3] People who are sensitive about something should take measures to conceal it from others in order to keep it private (e.g., keeping it in places culturally designated for private things). If they do not take these measures, others cannot be accused of wrongdoing when accessing the object of sensitivity.

- it is the kind of information that, if others knew about it, would put *S* in serious danger of being harmed.

On my account, *I* accesses *p*(i) the moment *I* knows about *p*(i). In other words, it is not enough for *p*(i) to be *accessible* to *I* for *S* to lose privacy. If *I* has *S*'s diary (i.e., it is accessible to him) but never reads it, *S* has lost no privacy. As long as *S*'s diary remains *unaccessed, S*'s privacy remains intact.[4]

Autotopos[5] refers to a metaphorical sensorial personal space, a sensorial self-space. In more colloquial terms, the autotopos is a personal zone of privacy that surrounds us wherever we go and sets the limit to others' accessing us sensorially. Its borders are set by cultural norms. It is the kind of sensorial space that people in *S* commonly would not want anyone, other than themselves (and perhaps a very limited number of other people chosen by them), to access.

There are two ways *S*'s autotopos can be accessed: 1) *S*'s autotopos is accessed when *I* sensorially enters a culturally established personal zone of his. That is, when *I* (through direct or indirect perception such as cameras and microphones) sees, hears, smells, or touches him in a zone where there are cultural expectations to be free from the eyes, ears, touch, and presence of others (e.g., in the toilet). 2)

[4] I realise that the adjective 'unaccessed' is not found in any dictionary, but there is no suitable existing term to convey in one word the property of not having been accessed. 'Inaccessible' denotes the property of not being able to be accessed, which is different from being accessible yet not actually accessed. Analogous differentiations exist in English, however, that use the same prefixes (e.g., indisputable/undisputed, inalterable/unaltered, etc.).

[5] From the Greek *auto* (self) and *topos* (place or space). My thanks to Roger Crisp for suggesting this term.

S's autotopos is also accessed when *S* is witnessed engaging in some activity or being the subject of some event that typically evokes the desire to have no witnesses or very few chosen witnesses (e.g., being naked). Such an activity or event may happen in the public sphere, as we saw in Chapter Two with the example of a person needing medical attention on the street.

One might wonder what makes the two species—informational and sensory access—part of the genus of privacy. The unity of the category of privacy is founded on the notion of being personally unaccessed and the kinds of interests we have in not being accessed by others. As I indicated in Chapter One, privacy evolved from a notion of territoriality or personal space that became metaphorical with the coming of language, and it protects us from a) certain kinds of harms that may come about as a result of other people having access to our personal life, b) the demands of sociality, c) being judged and possibly ridiculed by others (and thus from self-conscious negative emotions such as shame and embarrassment), and d) the discomfort of being watched, heard, and so on.

Some further clarifications are in order.

The autotopos is not a physical space

The autotopos is metaphorical because it is not a physical or territorial space. It is a *sensorial* space: a zone where one does not want to be sensed by others. Someone placing a camera outside of one's bedroom window may not invade a physical space, but it does invade a sensorial space. Someone placing a foot on a land that

one has never visited and yet owns halfway across the world may be an invasion of private property, but not of one's privacy. For privacy, what matters is that one can be seen, heard, smelled, and touched by others.

Privacy is not like a bar of chocolate

Talk of loss of privacy may give the impression that privacy is a set good, like a bar of chocolate or a pie, that gets diminished every time we lose some privacy, so that we may at some point run out of it. This interpretation is incorrect for two reasons.

First, we are always generating personal information and opportunities for sensorial access. Second, given that privacy can only be lost with respect to some information or sensorial access *p* and some individual *I*, and given that there are billions of people in the world, in practice one can always lose privacy to someone new. Having no privacy at all would mean being accessed (with regard to *all* personal information and one's autotopos) by all 7 billion people in the world at all times.

That said, one can say that one person has less privacy than another with respect to something specific, such as health information. If one person's medical records were leaked to thousands of people and another's medical records were not, if other things are equal, the first person has less privacy than the second with respect to health information.

Privacy does not track peace of mind

It is worth noting that the *condition* of having privacy (or not) and the *psychological state* of feeling secure in one's privacy may come apart.

Enjoying peace of mind as a result of believing one has privacy with respect to some *p* does not guarantee that one in fact has privacy with respect to *p*. Before the Snowden revelations, most people probably thought they enjoyed privacy with respect to their emails and Internet searches. They were wrong. The revelation that we have been watched may bring about embarrassment and anxiety, but those feelings do not mean that we have lost more privacy than when we were unaware of the spying.

Conversely, feeling anxiety as a result of believing one does not have privacy with respect to some *p* does not mean that one does not have privacy with respect to *p*. Employees may feel surveilled by a camera that is looking towards them, even if the camera is a sham one that is not actually recording what they do. Employees are not losing privacy here, even if they are made to feel and act as if they were.

It is wrong to mislead people about whether they have privacy or not, but this wrong is not the same wrong as that of invading people's privacy. I will come back to these kinds of cases in Section VII.

Some objections and responses

Before I go on to define the right to privacy, I will respond to some possible objections to my definition of privacy. My account relies on the notion of access:

informational access and sensorial access. Julie Inness (1992), whose own account of privacy I rejected in Chapter Three, offers the most challenging objections I have encountered to access-based definitions of privacy.[6] She thinks that even if 'privacy often *manifests* itself in conjunction with limited access or separation,' that 'does not mean that privacy *is* them' (43). Inness believes that privacy cannot be access-based for three reasons:

Objection 1. It would imply privacy has a neutral valence, 'since separation [as in, separation of others] is a neutral concept until it is placed within a particular context' (43). Yet she argues privacy is clearly positively valued, as ordinary language reveals in phrases such as 'enjoying privacy' and 'invasion of privacy' (44). Inness points out that

> …a number of words describe conditions in which the individual and her activities are forcibly separated from the access of others, for example, censorship, isolation, deprivation, but these words would seldom be replaced with "privacy" and retain the appropriate undesirable value connotation. (…)
> Saying that privacy, liberty, or free speech are positively valued states is equivalent to claiming that the burden of proof is upon the person who wishes to curtail privacy, liberty, or free speech—that person must offer a justification of her actions. This is not the case with a negatively valued concept such as "isolation": no justification is required for rescuing someone from isolation. (44-45)

It is questionable that one does not need a justification to 'rescue' someone from isolation. Hermits seek isolation, and it seems to me one would need a good justification to drag them back into a community. There is some truth in Inness's

[6] Kevin Macnish (2016) has very similar criticisms to my own of Inness' objections.

point, however: privacy is indeed generally valued positively. Inness is wrong, however, to think that the positive connotations often associated with privacy make it the case that privacy is not a matter of limiting access.

Privacy can be positively valued without being a good thing *absolutely*. In other words, that privacy is generally a desirable thing does not mean that the more privacy you have, the better it is in every conceivable circumstance. Speaking freely is a good thing; that does not mean that one should speak one's mind in every circumstance, at all times. Being free to leave one's home whenever one wants is a good thing; it does not follow that one should never go back home to get a good night's rest. Rather, one should have the ability to leave home and speak freely whenever one finds it good and desirable to do so. In the same way, privacy is a good thing in the sense that it is a good thing to be able to have whenever one finds it good and desirable to have it, or whenever it benefits our wellbeing and that of others on account of the kinds of creatures we are and the kinds of societies we have or want to have. Limits to privacy are desirable, lest we want to become completely isolated from each other.

It is perfectly possible to have too much privacy. Someone who is too private is someone who is not sharing his intimate thoughts with anyone, and is therefore missing out on the special kind of connections one can only establish through surrendering some of one's privacy and becoming vulnerable to another.

Objection 2. Thinking of privacy as separation from others makes privacy 'necessarily individualistic. As soon as one individual encounters another, no matter the nature of the encounter, privacy is necessarily lost' (43-44).

Inness's first mistake is thinking of privacy in an all-encompassing and binary way, as something you either have or lack. As has been mentioned, privacy should be thought of as something one has or does not have with respect to some aspect (some bit of information or some kind of personal access) and to someone. Consider the case of a couple having an honest and quiet conversation in their home about whether to have children. While neither spouse has privacy with respect to the other and with regard to their feelings about having children (assuming they are being honest), the couple have privacy with respect to the passers-by who cannot hear or see them.

Inness is sceptical that we lose privacy when we invite a close friend to our home, when we have consensual sexual intercourse with a lover, or when we allow a trusted friend to read a personal letter. She believes that it goes against our linguistic and moral intuitions to claim there is a loss of privacy in those situations: 'Consider telling the other people involved in these examples, "I appreciate your lessening of my privacy." (…) Our impulse in these cases is to say that we are including another within our realm of privacy, not lessening our privacy (even in a desirable fashion)' (46).

I disagree about this being the right way to think or talk about privacy. Inness' observation regarding linguistic usage does not prove as much as she thinks it does. Of course we would not thank our close ones for lessening our privacy, but we could admit (at least to ourselves) that surrendering our privacy to loved ones can be hard, even when it is worth it. Imagine a person who is insecure about her body because she has a scar, and thus feels fearful of losing the privacy in her nakedness with her lover. She might be, overall, happy to do it in the hopes that he will react well and not care about her scar, but she can still feel the apprehension of losing privacy to another.

When we lose privacy we become vulnerable, and it is in virtue of this disclosed vulnerability that surrendering privacy to someone is a gesture of trust that, when received with sensitivity, strengthens relationships. In a similar fashion, it might be hard to share with a friend a letter in which intimate details about ourselves are revealed, but we might be happy to do it in exchange for our friend's understanding, acceptance, and perhaps advice.

It is misguided to suppose that, because we invite another into our life (e.g., a new romantic partner), we automatically include her within our realm of privacy without losing any privacy. It would mean that romantic partners could never lose privacy with respect to each other. Someone who has her emails read by a jealous spouse would surely disagree.[7] It would also imply that one could not lose privacy

[7] At least in some countries, the law recognises that spouses can lose privacy to one another and violate one another's right to privacy. In Spain, spying on one's spouse's mobile phone is punishable with two years of prison (Precedo 2015).

through Facebook by sharing personal information, as one would merely be inviting one's Facebook friends into one's realm of privacy and thereby expanding that realm. I find these implications too counterintuitive to accept.

Objection 3. Innes argues that privacy cannot be access-based because many 'true privacy violations' become only 'threatened privacy violations' if one conceptualises privacy as a matter of being unaccessed (46). To illustrate her point, Inness suggests two examples:

> (1) I realize that a peeping Tom is coming to my window, so I evade him by ducking under the bed; (2) a home dweller realizes that someone is attempting to overhear her conversation (but has not done so), so she drags her friend into the closet to continue the conversation. (Inness 1992, 46)

Inness suggests that these cases involve a violation of privacy, and not just a threat. Access-based accounts of privacy, argues Inness, would be compelled to claim that as soon as the Peeping Tom manages to catch a glimpse of my foot or the eavesdropper manages to catch one word of the conversation, then the situation is transformed from a threat into a privacy violation despite the triviality of what is seen or heard. She believes that there is no salient difference between the Peeping Tom coming to my window and not seeing me, and the Peeping Tom coming to my window and seeing my foot. (A problem for Inness is that it is not clear she can distinguish between the moral difference between Tom seeing my foot and Tom seeing much more than that because both are simply privacy violations.)

Inness's third objection is neutralised when we specify that access (informational or sensorial) must be of a *personal* kind. If the eavesdropper catches only one word of a conversation (say, 'the'), or the Peeping Tom catches a glimpse of my foot or my hand, then that cannot possibly be the kind of access and information that people in a society would not want widely shared.

However, it would be understandable for a reader to think that my account does fall prey to the objection that many apparent *invasions* of privacy (in my terminology) are only threatened invasions. So far, my definition cannot account for the wrongness involved in the Peeping Tom's *attempt* to invade my privacy. According to my account there can only be an invasion of privacy when information is *known* (not possessed in other ways), when someone is *seen* (not when images of her are kept but not seen). This is problematic. It also implies that when the NSA collects data from people—emails, images from their webcams, recordings from their conversations on Skype—there is no loss (and therefore no invasion) of privacy. It is not until an analyst reads the emails, watches the footage, or listens to the conversations, that privacy is invaded and a loss of privacy is incurred.

I will argue that both the wrongness of attempting to invade privacy and that of collecting data without accessing it are not invasions of privacy, but rather violations of our right to privacy. Before I go on to present my account of the right to privacy, I will first explore the value of privacy, as I assume that rights are grounded in interests understood to be aspects of people's wellbeing.

IV. The value of privacy

Privacy offers many benefits. Some of them are individual, others are social, and yet others are political. The goods that privacy can deliver can also be categorised into those that are (partly) achieved through privacy, but could be achieved through other means, and those that can be procured only through privacy.

Financial and physical security are among the benefits that can be achieved in more than one way. We keep our credit card numbers from the eyes of others because we want to avoid identity theft. Depending on the country we live in or the job we have, we may be adamant in keeping our home address private in order to minimise risks of physical assault or kidnap. These benefits, however, could in principle be achieved through other means, including effective law enforcement. Using biometrics in addition to credit card numbers could solve the first problem without resorting to privacy. Hiring bodyguards could take care of the second risk. (Both measures would, however, erode privacy.)

While some of the benefits that privacy can offer may be achieved through other means, in the world as it is today, ensuring privacy is very often the most cost-effective way to avoid certain harms people might suffer if they become personally accessed. The cost-effectiveness of privacy, however, is highly contingent and may change with the advent of new technologies and social developments.

Paramount among the goods in the second category—benefits that only privacy can deliver—is 'people's interest in having a reasonable measure of control over the ways in which they can present themselves (and what is theirs) to others' (Marmor 2015, 3-4).

In his book, *The Presentation of the Self in Everyday Life* (1959), sociologist Erving Goffman described how this control was crucial for successfully managing both professional and personal relationships. Similarly, James Rachels (1975) has pointed out that being able to control who has access to oneself and information about one is intimately related to one's ability to maintain different kinds of relationships. One shares different aspects of oneself with one's students, one's best friend, and one's spouse.

It could be objected that we should not be given the chance to act differently in the presence of different people because that amounts to being two-faced, dishonest, or inauthentic. On this view, a complete lack of privacy (radical transparency) would be good, because it would force people to show their true nature to everyone alike.

As Rachels has argued, however, different relationships are partly defined by different patterns of behaviour (1975, 326). One's best friend would not be one's best friend if one could not cry, swear, and express one's fears in her presence. Displaying the same kind of behaviour with one's students, however, may be inappropriate. It is not that one is being dishonest when one acts with students

differently from the way one does in the company of friends. It is more that relationships function as a kind of division of labour, and burdening one kind of relationship with the load of another type of relationship creates confusion and dissatisfaction. Students typically do not want to hear about one's personal troubles—only good friends do (or should). Students want (or should want) to hear about the topic one is teaching. Furthermore, as Thomas Nagel has pointed out, these patterns of behaviour are not dishonest because they are not meant to deceive; they are social conventions that are typically shared and well known by all (1998, 7-8, 11). Being able to conceal certain aspects of oneself in order to present oneself in appropriate ways depending on the public one is facing 'protects one from the sense of exposure without having to be in any way dishonest or deceptive, just as clothing does not conceal the fact that one is naked underneath' (Nagel 1998, 8).

Even if one concedes that privacy can and does facilitate one's being dishonest, it is important to remember that there are limits to the value of honesty. Being forced to speak our minds fully and sincerely on every occasion and with every person we meet would be disastrous. Conventions of restraint prevent unnecessary conflict and complications (Nagel 1998, 9). If there were no concealment, if all our political views, religious beliefs, sexual practices, and more, were exposed at any place and time, there would be more opportunities for unnecessary public confrontations about issues that could well remain private without any losses, for the benefit of all.

Having some measure of control over how we present ourselves to others also protects us from being judged on aspects of ourselves we do not wish to share with others—it protects our reputation, it saves us from embarrassment and prejudices, and it creates a safe space to experiment and unwind.

Reputation is important for people to thrive in social settings. In a world without privacy, a Jean Valjean could never become Monsieur Madeleine. Sometimes hiding our past, our heritage, or some other information about oneself is the only way to transform one's life for the better, to avoid unfair discrimination, and escape having others categorise one into an unmovable box.

A major motivation to protect privacy is avoiding embarrassment and other negative self-conscious emotions. There are certain morally innocuous things that we can only do at ease by ourselves or with select intimate others. The unwanted presence of others can either eradicate the possibility of doing these things, or transform the nature of activities in undesirable ways (e.g., what could be a meaningful and intimate conversation with one's spouse may become a superficial and bland one in the presence of another).

Privacy enables people to practice new skills without fear of being ridiculed. It also allows people to experiment and try out ideas they may not wish to endorse in public after having thought them through in private. Privacy promotes autonomy, independent thinking, and creativity by providing people with a space where they

can be free from pressures to conform. It allows individuals to do what they would not do in public out of fear of what others might think of them.

As we explored in Chapter One, a crucial interest we have in enjoying privacy is obtaining respites from the burdens of sociality. Having some time and space free from others' gaze contributes to people's psychological wellbeing. Engaging in social interaction means having to fulfil expectations and deal with responsibilities; keeping others at some distance from time to time allows us to relax. Privacy also frees us from unwanted distractions; it gives us opportunities to concentrate (Gavison 1980, 446-447). Having a zone where we can be free of intrusions enables us to devote our time and energy as we see fit without having to worry about staging a performance for others or catering to their needs.

In the political realm, privacy contributes to the protection of liberal, democratic, and pluralistic societies (Gavison 1980, 442). It fosters and encourages individuals' autonomy by shielding them from external interference. Political liberty requires that people have the right to keep private their votes, associations, and thoughts, if they so wish. As Gavison puts it, '[p]rivacy is crucial to democracy in providing the opportunity for parties to work out their political positions, and to compromise with opposing factions, before subjecting their positions to public scrutiny' (1980, 456).

When democracies erode and disagreeing with power becomes dangerous, privacy protects political dissenters such as activists and whistleblowers. When advocating

against injustice may endanger one's life or one's loved ones, having the possibility of anonymous protest and resistance becomes crucial to defending democratic ideals.

The right to privacy rests on the importance of the interests people have in (thin and robust) privacy. The importance we attribute to the protection of those interests results from their contribution to a public culture that enhances individuals' wellbeing (Raz 1988, 256). The importance of the right to privacy, like other liberal rights, lies both in its service to individuals and the public good. The following section explores what it means to have a right to privacy.

V. The right to privacy

The right to privacy as a right to a robustly demanding good

Let us start with an example. Mercer notices that Mae forgot her diary at his house. He decides not to read it, and gives it back to her the next day. I suspect most people would agree that Mae has lost no privacy in this case. Mercer could have read her diary (thereby making her lose her privacy), but he did not.

Now compare this case with that of the NSA, which was left pending a section before last. Suppose the NSA has collected all of Mae's emails, footage from her webcam, and all her Skype conversations. Suppose further that nobody has ever read those emails, looked at the footage or heard the conversations. Are we not

committed to say that, if in the first case no privacy is lost, no privacy has been lost in the latter case?

It might be objected that the diary case and the NSA case are not analogous, because in the diary case Mercer did not *intentionally* come to possess Mae's diary, whereas the NSA intentionally gathers our data in case they might want to look at it in the future. It is as if the NSA had *stolen* our diary. But surely the difference in the intention of the person or entity that possesses our information creates a difference in the moral status of *their* action (the kind of action it is, the wrongness of it) and *not* a difference in the privacy Mae loses. Mae does not lose privacy in either case, because no one has read her diary or her emails; no one has learned anything intimate or sensitive about her. In the NSA case, however, it is plausible to think that Mae is wronged in a way in which she is not wronged in the diary case.

The wrong that Mae suffers signals that, when people demand that their 'right to privacy' be respected, they are not only wishing not to lose their privacy—there is also an implicit demand for *robust* privacy. In everyday speech, when we talk of privacy, sometimes we refer to thin privacy (as when we say that Mae does not lose privacy if Mercer does not read her diary) and sometimes we refer to robust privacy (as when people claim their privacy has been violated by the NSA even if no one has read their emails). My aim is to separate these two for the purposes of clarity and precision.

I contend that the good of privacy is a minimally demanding or actual one—it is one we either have in the here and now, or we do not. The right to privacy, on the other hand, is a right to a rich or robustly demanding good. Robustly demanding goods are ones that require counterfactual assurances. The right to privacy requires, not only that you not invade my privacy here and now, but that you would not invade my privacy in a range of relevant possible situations (e.g., if you stopped liking me or if invading my privacy suddenly became profitable for you).[8]

Thus, for Mercer to respect Mae's right to privacy, he not only must refrain from reading Mae's diary if she forgets it in his home. He must also be disposed not to read it without her consent even if he did not like Mae, or even if he was very curious to see if Mae had written about him, or even if he could stand to benefit from profitable information in her diary, etc. If Mercer is not unwilling to read the diary, if he does not have a disposition and commitment to refrain from reading Mae's diary but he does not get a chance to read it (e.g., because Mae does not forget it in his house), he is not violating Mae's right to privacy, but he is failing to respect it.

If Mercer secures a position from which he can read Mae's diary, with the intention of reading it now, in the future, or in some counterfactual situation (e.g.,

[8] Rich goods have a structure that mirrors the Republican ideal of freedom. For Republicans, it is not enough for someone not to suffer actual interferences to be free. A slave might have a master that has never interfered with him and still not be free. As long as someone could interfere with one arbitrarily (i.e., with impunity), one is not free (Pettit 1996). Pettit (2015) has used this structure to argue that goods such as love, virtue, and respect are also counterfactually demanding in this way. I wish to include the right to privacy in this list.

he steals the keys to the drawer where she keeps her diary with the intention of reading it at the first opportunity he gets), attempts to read the diary (e.g., he tries to read it but cannot understand Mae's handwriting), or succeeds in reading her diary, he violates her right to privacy. The third option is morally worse than the second, and the second is morally worse than the first. The latter option is the worst because it includes a rights violation *and* a loss of privacy (the worse the loss of privacy, the morally worse the violation is). The second option is worse than the first because it puts the victim's privacy at greater risk.

More generally, *I* fails to respect *S*'s right to privacy despite respecting *S*'s privacy here and now if *I* does not do so robustly: if *I* has not invaded *S*'s privacy out of luck or laziness or if *I* would be ready to invade *S*'s privacy in relevant possible worlds. Furthermore, absent outweighing conflicting considerations, *I* violates or infringes (if the act is justified) *S*'s right to privacy if he

a) secures a position from which he can invade *S*'s privacy with the intention of invading *S*'s privacy now, in the future, or in some counterfactual circumstance, or
b) attempts to invade *S*'s privacy, or
c) invades *S*'s privacy.

This account includes cases where a person gains access to someone's privacy with his consent, but then betrays his trust and exposes his privacy to others. Suppose Mae shares her diary with Mercer but asks him to keep it to himself. Mercer decides to betray her and publishes her diary in order to hurt her. Given that Mae

is being made to lose privacy without her consent, this example would count as an unjustified invasion of privacy (c), and therefore a rights violation.

Where conflicting considerations outweigh the interests of the would-be-right-holder, it is not justified to hold people subject to duties to protect the privacy of the would-be-right-holder (Raz 1984, 210); the right to privacy would then be infringed, but not violated, as an infringement signals a justified invasion of privacy. Suppose that Mae is not around to ask for her consent, but the police are confident that in her diary is information to save a thousand lives. There is no other way to save these lives. In this case, Mae's right to privacy is infringed but not violated if the police use her diary to save those lives.

Now that the right to privacy has been somewhat clarified, we can go back to the NSA case. To put it tentatively, then, the NSA violates our right to privacy by collecting our personal information because they put themselves in a position from which they can invade our privacy (i.e., make us lose privacy without our consent) at any moment, in many possible worlds, even if, in the actual world, they never do, and we never lose privacy. If an analyst from the NSA does read our emails, then our right to privacy is also violated, to an even graver degree, because we also lose privacy. Our right to privacy asks of the NSA that it refrain from collecting and accessing our private data in relevant possible worlds. It might be objected that the NSA does not violate our right to privacy, but merely infringes it, because there *are* outweighing conflicting considerations (i.e., security

considerations that are more important than privacy). This is the topic of Chapter Five.

The lack of connection between the right to privacy and privacy—the fact that one's right to privacy may be violated even if one does not lose privacy—justifies my claim at the very beginning of this chapter that the right to privacy is not concerned with the object of privacy (private information or sensorial access), but rather with counterfactual circumstances and with ways of getting at objects of privacy (i.e., bypassing consent).[9]

Again, a general theory of rights is beyond the scope of this chapter, but it seems to me that at least some rights—those that protect robust goods—share the same structure. Democracy and security are robustly demanding goods (Southwood 2015, Lazar 2015), and the rights to democracy and security share some characteristics with the right to privacy.

Take the right to democracy. An imperialist occupying polity fails to respect the occupied polity's right to democracy if it does not do so robustly. It is not enough that it allows the occupied polity to carry on as before being occupied. For it to

[9] Andrei Marmor (2015) has also argued that the right to privacy is about ways in which information is obtained. He believes, however, that A violates B's right to privacy only when A manipulates B's environment in a way that either diminishes B's ability to control information about herself or reduces the options B can choose from. Marmor's view cannot account for cases such as that of Mae forgetting her diary at Mercer's house, because he did not manipulate Mae's environment. It seems clear, however, that Mercer's reading Mae's diary simply because he happens to have the chance is a violation of her right to privacy. Similarly, Marmor cannot account for the moral difference between the NSA *collecting* data and an analyst *accessing* the data.

respect the right to democracy of the occupied polity, it is also necessary that it not be willing to intervene across relevant possible worlds—across certain changes in the content of the will of the occupiers and the occupied (Southwood 2015). The occupying force violates the right to democracy of the occupied if it secures a position from which it can intervene with the intention of intervening (e.g., places snipers at the top of buildings who are ready to fire if the occupied act in an undesired way), attempts to intervene (e.g., the snipers shoot but miss), and if it successfully intervenes with the will of the occupied people (e.g., the snipers kill or hurt occupiers). Here too, the third option is morally worse than the second, and the second is morally worse than the first.

Those who are not sure about the label may call *robust privacy* (as opposed to thin or actual privacy) what I am calling the right to privacy. One reason to call it a right, however, is that it grounds imposing requirements on others' behaviour. Robust privacy grounds duties: it asks people to show restraint and refrain from invading others' privacy, absent outweighing conflicting considerations. Requiring certain actions from others is characteristic of rights. In the words of Joseph Raz: 'Rights ground requirements for action in the interest of other beings'; 'the special features of rights are their source in individual interest and their peremptory force' (1988, 180, 192).

An objector might want to be more precise and talk about a right against securing a position from which someone can invade one's robust privacy, and a right against attempts to invade one's robust privacy, as well as a right against invasions

of one's robust privacy. For the sake of brevity, however, I will continue to talk about the right to privacy.

I do not endorse strict correlativity between Hohfeldian incidents (Hohfeld 1919), but his terms are helpful in clarifying the nature of rights. Although Hohfeld wrote about legal relations, his analytical scheme can apply to moral relationships (Kramer 2000, 8). In Hohfeldian terms, the right to privacy is first and foremost a claim-right, as it grounds duties in others. For others to respect our right to privacy, they must have a robust unwillingness to invade our privacy.

The right to privacy also encompasses an immunity, as others lack the (moral) ability to alter our claim to privacy. Our immunity prevents others from waiving or annulling our claim to privacy. For example, others cannot waive or annul our right that information about our sexuality be kept private. If other people have information about our sexuality, they should not disclose it. In this sense, it is a passive right, as it regulates the actions of others, and not of the right-holder.

There is also an active side to the right to privacy, however, in the sense that it refers to the actions of the agent, as it is likewise a privilege or a liberty. Our right to keep our privacy means we do not have a duty not to keep our personal information and autotopos unaccessed. Therefore, we are at liberty not to disclose our personal information and not to allow access to our autotopos.

The right to privacy, then, is primarily a claim-right amounting to the assurance that other people will refrain from invading our privacy (i.e., accessing our private information or autotopos without consent) in relevant possible worlds. I will go on to detail what the demands of the right to privacy cover—the range of possible worlds over which one is required to show restraint towards someone's privacy in order to respect that person's right to privacy—and what is meant by 'assurance.'

The counterfactual demands of the right to privacy

Enjoying a right to privacy means that others respect our privacy in the actual world and in relevant possible situations. There must be some limit, however, to the range of possible worlds included for this requirement to be met. For a possible world to count as relevant, circumstances must be such that moral reasons to respect the right to privacy continue to outweigh the balance of competing considerations. This constraint excludes cases where there are reasons to invade another's privacy, but these reasons are either not moral in nature (i.e., curiosity, personal profit) or they are moral reasons that do not outweigh the right to privacy.

The right to privacy grounds a duty on the part of others to refrain from invading our privacy. If there are outweighing conflicting considerations, however, the right to privacy cannot justify those duties. One respects *S*'s right to privacy provided one has the disposition and commitment that one will not invade *S*'s privacy as long as there are no moral reasons that outweigh that right. When moral reasons

heavily outweigh someone's claim to privacy, privacy can be justifiably invaded—with appropriate compensation to the victim if she is innocent.

Let us suppose the police have strong reason to think Raskolnikov is guilty of a serious crime. They investigate him and in the course of the investigation, justifiably invade his privacy by wiretapping his phone for a week. The invasion is justified because we are supposing the police have good reasons to believe he is guilty of a serious crime, and ascertaining culpability and protecting the public from a criminal repeating a crime outweigh the amount of privacy Raskolnikov loses. No right is violated; the right to privacy is merely infringed. But if he turns out to be innocent, compensation is owed to him on account of his having lost his privacy through no fault of his own.

I now go on to detail what I mean by 'assurances' when I say that the right to privacy amounts to having an assurance that others will not invade our privacy.

Assurances

My account of the right to privacy has a built-in rule-of-law requirement. In order to enjoy assurances that others (government officials and private citizens) do not invade our privacy, now or in the future, laws must be in place to prevent that from happening and, if it happens, to punish offenders. In Christian List's words:

> Formally, these considerations suggest that to achieve the rule of law, we must organize the world (through institutional design or policy interventions) in such a way that whenever a normative

> law is given in the form of a modal desideratum **O***P* [It is obligatory that *P*], then a corresponding positive law holds in the form of a modal fact □*P* [It is necessary that *P*]. Such a modal fact will typically not be as robust as a law of physics: *P* will be true neither in all physically possible worlds nor even in all socially possible worlds in a broad sense. But, ideally, *P* will be true in a large range of socially possible worlds relative to the appropriate social background conditions—those shaped by our institutional design or policy interventions. (List 2006, 209)

List writes about Republican freedom, but the same point applies to the right to privacy.[10] Institutional arrangements and laws can make good on the moral entitlement that others respect our privacy across relevant possible worlds.

Let us return to the example of the NSA. It can be said that the intelligence agency violates our right to privacy partly because there are no legal or institutional measures to make sure that people's privacy will not be invaded in unjustified ways.[11] Consider the following excerpt from an interview given by Edward Snowden:

> Many of the people [intelligence analysts at the NSA] searching through the haystacks were young, enlisted guys and … 18 to 22 years old. They've suddenly been thrust into a position of extraordinary responsibility where they now have access to all your private records. In the course of their daily work they stumble across something that is completely unrelated to their work, for example an intimate nude photo of someone in a sexually compromising situation but they're extremely attractive. So what do they do? They turn around in their chair and they

[10] Recently there have been at least two attempts to offer a Republican account of privacy (Roberts 2015, Newell 2014). Both argue that the value of privacy lies in its usefulness to protect citizens from domination. They do not suggest privacy is a robust good.

[11] Even if there were legal constraints in place, the NSA might still be accused of violating people's privacy if they are ready to invade privacy in circumstances where the moral reasons to respect the right to privacy continue to outweigh the balance of competing considerations. Chapter Five deals with this issue in depth.

> show a co-worker. And their co-worker says: "Oh, hey, that's great. Send that to Bill down the way." And then Bill sends it to George, George sends it to Tom and sooner or later this person's whole life has been seen by all of these other people. (…) It's never reported, nobody ever knows about it, because the auditing of these systems is incredibly weak. (Snowden 2014)

As Snowden's testimony shows, there are no policy mechanisms—no institutional arrangements, no system of oversight, and no laws—to punish and prevent abuse of this kind from NSA analysts.

One might think that the kind of cases Snowden brings up do not count as abuse because the people being surveilled were likely criminals. Criminals who have committed a grave offence can justifiably have their right to privacy infringed because there are the outweighing conflicting considerations of ascertaining their culpability (once there is strong suspicion they are criminals) and ensuring public safety (preventing serious crimes). One can argue, however, that criminals' privacy should be invaded only as a means to either of those goals, but not for entertainment, ridicule, or punishment. Any invasion of privacy beyond those outweighing conflicting considerations would be a rights violation. It is beyond the scope of this paper to determine what is necessary, in terms of privacy invasions, to ascertain culpability and prevent serious crimes, but it seems safe to say that in the vast majority of cases, it will not include spying on people's sexual lives (sex crimes apart) and sharing pictures with colleagues of attractive naked people just for the fun of it.

Furthermore, the NSA does not spy only on people who are criminals or suspected criminals. The NSA aspires to 'collect it all' (i.e., collect all possible data

from everyone in the world) (Greenwald 2013). The former head of the NSA, General Michael Hayden, admitted that '[t]his is not about guilt… NSA doesn't just listen to bad people. NSA listens to interesting people' (cited by Friedersdorf 2015). The NSA violates people's right to privacy because it is ready to invade people's privacy whenever it can, with no discrimination as to how it collects data, or what kind of information is being accessed, and with few legal limits, as we will further see in Chapter Five.

Because of the built-in rule-of-law requirement, the moral and the legal right to privacy are intimately intertwined. That means that when the legal right to privacy is violated, the moral right to privacy is violated as well. The converse, however, is likely false. The realm of the legal seems narrower than the realm of the moral, as not all violations of the right to privacy seem to merit legal action (e.g., consider a friend taking a peek at another's friend's passport to ascertain his birthday). This topic, however, is too broad to take on here.

Laws and institutional mechanisms of oversight are not the only type of assurances we can enjoy. When it comes to personal relationships (i.e., familial relationships or friendships), other people's disposition and sincere commitment to respect our privacy are also guarantees (even if they are not perfectly robust and somewhat fallible). When a spouse enjoys her partner's respect of her right to privacy it is not only because (or even primarily because) her partner is subject to laws that incentivise him to respect his spouse's right to privacy, but also because he is committed and has the appropriate disposition to respect her privacy now, in the

future, and in a variety of possible worlds (including worlds where he would not be caught by the police, or where there were no laws against snooping). Commitment to respecting privacy is important because it is unsatisfying to have someone not invade one's privacy merely because she is not interested enough, or merely because she is subject to laws—the moment something in one's private life should become interesting to her, or the moment she could get away with it, she would invade one's privacy.

The robust disposition to respect our privacy in relevant possible worlds is likewise important because commitment is not enough. Some people are notoriously bad at making good on their commitments. People's capacity to stick to their commitments is important to be assured that we will respect each other's right to privacy.

Commitments and dispositions to privacy are not only characteristics of individuals. They can also be characteristics of the culture in a society; they can be more or less entrenched in social norms. Since social norms have come up in the definition of privacy, as well as in the sections dedicated to the value of privacy and the right to privacy, it is worth looking at what role they play.

VI. The role of social norms

Social norms are important for privacy and the right to privacy. They are important in partly determining what is to be kept private in a culture, and in disciplining social interactions in a way that protects our privacy.

Social norms are general tendencies people have to behave according to conventions that are widely shared in a society. When a convention is a social norm, there is a general expectation that people in that society will approve of acts that respect the convention and disapprove of those acts that do not accord with it (Pettit 2015, 37). If a member of a society were to publicly defend a social norm, other people might be in disagreement, but no one would doubt such a norm, as members of a society recognise and are familiar with the social norms that rule their lives (Pettit 2015, 39). When someone breaks the social norm, people around her are likely to signal disapproval.

Privacy norms are among the clearest instances of social norms. If you break them, people around you will likely let you know with anything from an awkward silence, to an angry demand for an explanation, or an insult. Anyone having doubts about the strength of social norms about privacy can try standing behind a colleague's computer at the office while she reads her emails and reading her messages out loud. Better yet, take a peek over the toilet stall wall in a public toilet

and greet the busy occupant. You will not be disappointed—privacy norms are alive and well.

The desire for good reputation and social acceptance serves an important role in motivating people to acquiesce with privacy norms. Given that there is almost always a risk of getting caught, the wish to avoid social rejection usually outweighs any temptation people might have to invade others' privacy. Complying with privacy norms is further incentivised by the interest that every person has that her own right to privacy be respected.

Social norms are vital to ensure that privacy is protected because they inspire, complement, and support laws. A high reliance on social norms has the benefit of making people aware of the responsibility they have in cultivating the kinds of social norms that can support a culture that fosters wellbeing. Privacy social norms inform policy-makers. They also encourage compliance with privacy laws, create limits where internally-motivated dispositions and commitments fail, and many times are successful in preventing possible offenders from committing grave violations against the right to privacy that would need to be dealt with by the law.

A crucial objection to my account is that, by relying so much on social norms, I have lost the capacity to offer a purely descriptive definition of privacy—I have instead imported and rubberstamped social norms of privacy, and have also surrendered my ability to have a normative judgment on what should and should not be private. I disagree.

First, my definition of privacy is not normative, because it only *describes* social norms, it does not endorse them. As a result, it can be filled in with different substantial conceptions of what people value as private. It could be counter-argued that by giving importance to personal information and personal sensorial access, I am already endorsing the social norms of my own society, even if I do not specify what is to be considered personal information and personal sensorial access. As we saw in Chapter One, however, concern about information that is considered personal and about personal access of different kinds is universal across time and cultures.

Second, while it is true that my definitions of privacy and the right to privacy do not give me enough grounds to take a stand on what should be the substantive contents of privacy (beyond saying that privacy is about information and sensorial access about which people are sensitive), my view on the value of privacy does. Privacy is valuable to us as a tool that contributes to the attainment of certain objectives: safety, relaxation, independent thought, autonomy, political freedom, etc. In a nutshell, privacy is valuable because it can contribute to individuals' wellbeing. Let us suppose that in society *C*, people think that information *f* (people's sexual fantasies) should not be private, and should be published for everyone to see. Suppose, however, that having information *f* out in the open is counterproductive to people's wellbeing. It creates unnecessary jealousy, embarrassment, and conflicts. Hence, one can argue that *f* ought to be considered personal information, and therefore ought to be protected by the right to privacy.

The example used may be generalizable to many, most, or all societies. Other examples only hold for certain societies. Let us suppose that *f* stands for physical location, and let us further suppose, only for the sake of exegesis, that the only reason we have for caring about whether our location is disclosed is physical safety. In dangerous society *D*, *f* should be private because publishing that information can risk people's lives. In safe society *S*, *f* does not have to be private (or should not be private, if we find reasons that weigh in favour of making it public).

Consider a thornier example that involves more normative issues. Suppose that *f* stands for being gay. One might argue, then, that in homophobic society *H*, *f* should be private, because if that information becomes public, gays will face discrimination and might even have reason to fear for their lives (the implication being that they should not be outed by others). In society *J*, where homosexuality is fully accepted, *f* does not need to be private. The problem with this example is that, in order to transform a homophobic society into one that is not, it is reasonable to believe that it is necessary for homosexuals to 'come out of the closet,' even if that means making themselves temporarily vulnerable to discrimination and abuse. It is a sacrifice thought to be necessary for constructing a better society in the future.

The example of homosexuality suggests that, when it comes to issues that may not be generalizable to many societies, determining what should and should not be

private in a certain society on the basis of people's interests can inform us of how we should treat others (what to keep private about others), but sometimes it does not tell us enough to be able to decide what we should keep private about ourselves. When it comes to what we should keep private about ourselves, other considerations that may vary on a case-to-case basis come into play.

Because *H* is a homophobic society, our moral exploration tells us that it would be very wrong to publish *f* about someone, as we would be endangering his life. Whether or not one should make *f* public about oneself (supposing one is gay), however, will depend on a variety of issues, among which the following four are paramount: how risky it is for oneself in particular (the kinds of risks involved and the likelihood of incurring harms), whether one's sacrifice will contribute to transforming society for the better, whether one is willing to make that sacrifice, and whether there might be other reasons for keeping that information private. In cases of high risk, making something public for the purposes of bettering society should be considered a supererogatory act.

In the absence of an extremely weighty justification, the disadvantages and risks of losing privacy in hostile societies are too burdensome for some people to decide for others whether something kept private should be made public. That is a decision individuals must make for themselves.

VII. Perceptions of privacy and deceptive privacy

So far I have offered an account of privacy and of the right to privacy (as well as accounts of the value of privacy and the role of social norms for privacy). These explanations have not, however, covered all the possible harms and wrongs connected to privacy.

There are times when people think they have privacy but in fact are being surveilled. And there are other times when people think they are being surveilled when in fact they enjoy privacy. Sometimes these misperceptions happen unintentionally, as when a person is misinformed, say, about her employer's surveillance rules, even if these rules are publicly available, or when a person thinks someone is gazing at her, when in fact the person is blind.[12] Many times, however, people take the trouble to manipulate others' perceptions of privacy.

When someone succeeds in manipulating a person's perception of privacy so that she believes she has privacy when in fact she does not, it is a case of *deceptive privacy*. If she is made to believe she does not have privacy when in fact she does, it is a case of *deceptive lack of privacy*. These cases are notoriously absent from the literature

[12] In a *Just for Laughs* prank, a man stares intently at people who are sitting at a café. People get nervous quite soon: they choke on their food, they look back defiantly at the staring man, they turn their heads. Then the staring man puts on a pair of sunglasses and takes out his walking stick, acting as if he were blind. People relax in relief. (This particular man was not actually blind, as this was a prank, but we could suppose he was for the sake of illustration. He could have been blind, and the situation could have been real and not staged.) For a good laugh, see https://www.youtube.com/watch?v=hCpKNNtUwxA

on privacy, but they can be as morally significant as violations of the right to privacy.[13]

Consider a case of deceptive lack of privacy. Suppose O'Brien offers Winston the possibility of living in a nice apartment, free of rent, on the condition that he may install cameras in all rooms in order to surveil him at all times when he is at home. Compare two possible worlds. In the first world, the cameras are toy cameras that cannot record. O'Brien simply wants to make Winston believe he is being surveilled to make him behave well. In the second world, cameras are real and they do record. Which world is morally worse?

The first world has all the bad psychological and behavioural consequences of Winston perceiving his right to privacy being violated and his privacy being lost, and the wrong of Winston suffering deceit. Here, I take it that, even if there are no bad consequences resulting from deceit, other things equal, deceit is a wrong we should avoid if we can. The second world has all the psychological and behavioural consequences in Winston from the first world plus the wrongs, harms, and risks from the actual loss of privacy suffered by Winston (i.e., O'Brien having more information about him that he may use in the future to blackmail him, ridicule him, etc.).

[13] The only paper I have found in which privacy losses are clearly differentiated from privacy perceptions is Macnish's "Government Surveillance and Why Defining Privacy Matters in a Post-Snowden World" (2016), but the distinction is merely noted without any further exploration as to its moral significance.

In order to decide which world is worse, Winston being deceived must be weighed against the wrongs, harms, and risks of a violation of the right to privacy and the loss of privacy. Both deceit and the violation of the right to privacy are moral wrongs, and for the sake of simplification may be considered roughly equal wrongs. In the second world, however, we must take into account the added bad consequences and risks from the actual loss of privacy. Therefore, it is reasonable to think that the second world is morally worse—that it is worse for O'Brien to surveil Winston (with him noticing) than to fake surveil him. However, O'Brien is as responsible for the bad consequences in the psychology and behaviour of Winston in the first world than in the second. This last moral observation is analogous to Chicago's legislation on toy guns, according to which if a fake gun is used to commit a crime, the criminal is tried as if she had held a real gun.

Now consider a case of deceptive privacy. Terez and Tomas are a couple. Compare two possible worlds. In world one, Terez warns Tomas that she is a jealous spouse and will spy on him every chance she gets, even without his consent. In world two, Terez assures Tomas that she is committed to respecting his right to privacy. In fact, she spies on him as much as possible. Which world is morally worse?

World one has all the wrongs, harms, and risks of a violation of the right to privacy (assuming Tomas does not give his consent), and a loss of privacy, plus the bad psychological and behavioural consequences of Tomas perceiving his right to privacy being violated and his privacy being invaded. World two has all the

wrongs, harms, and risks of a violation of the right to privacy, and a loss of privacy, plus the wrong of deceit.

In order to decide which world is worse, one might be tempted to think that all it takes is balancing the bad psychological and behavioural consequences of Tomas perceiving his right to privacy being violated and his privacy being invaded in world one against the wrong of deceit in world two. It could be argued that these two roughly equate each other, or that one of them is weightier than the other. I will not express a view on this point. A further consideration is missing, however. In world two, Terez is likely to get much more information out of Tomas. Given that he does not think he is being spied on, he is likely to act naturally and freely, and as a result will suffer a much greater loss of privacy in world two than in world one (with all the bad consequences and risks that entails). For this reason I tend to think that world two is worse—that it is morally worse for Terez to spy on Tomas while deceiving him about it than to spy on him and tell him about it.

In real world scenarios, much will depend on empirical facts. Whether deceptive privacy or deceptive lack of privacy is worse than violations of the right of privacy will depend on the details of the case and is something that cannot be decided beforehand. Some of the relevant factors for a moral inquiry into a particular case include the victim's sensitivity to perceived violations of her right to privacy (the more sensitive the victim is, the more morally unacceptable it is to make her think she does not have privacy), how much privacy the victim will lose (in comparison to how much privacy she will lose if she knows she is being surveilled), and the

actual risks faced by the victim as a result of a loss of privacy (is the person violating her right to privacy likely to harm her using the information he gets from her?).

Even if the moral badness of real cases cannot be decided in advance, or even if one thinks that deceptions about privacy and lack of privacy are generally less morally grave than actual violations of the right to privacy, the upshot of this section is still valuable. The objective of this section has been to point out that some moral wrongs are shared in violations of the right to privacy, deceptive privacy, and deceptive lack of privacy; that we have the same kinds of interests in avoiding these wrongs; and that deception is morally significant and interacts with other morally significant factors in cases related to privacy (sometimes making the case morally worse, and other times making it morally better).

There is one final step to have a complete, yet rough, sketch of the moral map of privacy: obligations to keep something private.

VIII. Obligations to protect one's own and others' privacy

The discussion on the value of privacy prompted a defence for a right to privacy. The interests we have in privacy, along with the cases, distinctions, and arguments explored in this chapter, however, suggest that there is not only a right to privacy, but also certain privacy-related obligations and moral reasons to protect one's privacy and that of others.

One obligation, which might be better described as a duty (as it is grounded in the right to privacy) and has already been explored, is refraining from wrongfully invading other people's privacy. Beyond that general duty, there are moral reasons to protect our own and others' privacy. The following obligations or moral reasons to act should not be read as being correlative to rights. The world would be a better place if we complied with the following obligations because it would create a culture of privacy (with all its benefits), but it is not necessarily the case that the general interests in having a culture of privacy are enough to warrant rights beyond the right to privacy.

We have a right-based duty to respect privacy when there is an identifiable right-bearer; we are morally bound to respect her privacy because that is what is demanded of us by her right. We have privacy obligations when it is not clear who would be indirectly exposed by our indiscretion, and when exposing something is likely to harm one or more people in a way that could be described as wrongdoing on our part. We have moral reasons for action when failing to carry out the good act (protecting privacy) does not create enough of a moral loss to be considered wrongdoing. When we have obligations to protect privacy and when we merely have moral reasons to do so cannot be decided beforehand; it can only be assessed on a case-to-case basis. Much depends on context, and there are many grey areas where there may be much controversy and uncertainty about whether exposing something counts as wrongdoing or not.

Let us turn now to the obligation of protecting one's own privacy. When interacting with others or when engaging with websites on the Internet, we ought to be mindful of our privacy and show restraint regarding what we share about ourselves. By showing restraint I mean not exposing our privacy to just anyone or to a large group of people—none of what I say applies to what we should share or not with our loved ones or with trustworthy professionals like doctors and lawyers.

People who believe that one has duties to oneself, or self-regarding duties, might want to characterise some of these obligations (partly) as duties of self-care and/or self-respect (Allen 2011, 2013). On this view, we should not widely share embarrassing facts about ourselves or publish personal information that might put us at risk out of the duties of respect and care we owe to ourselves. Kantians can argue that self-regarding duties against the publication of sensitive facts about oneself are grounded in a duty to avoid actions that will undercut future opportunities to develop one's autonomy and freedom (Allen 2013, 857).

Self-regarding duties are controversial, however. Critics argue that they are nonsensical: having a duty to oneself would be like suing 'oneself in a court of law for return of the money one owes oneself' (Singer 1959, 202). Only the person to whom one owes a duty can release one from that duty. Owing a duty to oneself would mean that one could release oneself from the duty, so that the peremptory force of the duty would be lost.

One possible response to critics of self-regarding duties is to say that one's present self owes duties of care and respect to one's future self. It is unclear, however, when the present self ends and the future one begins. More importantly, if A (one's present self) is a different person from B (one's future self), B should not care (in a personal way) that A exposed personal information about herself, since A is a person who no longer exists and is different from B. There might be a plausible response here based on the historical and psychological connection that A and B share (and that is attributed to B by other people), but exploring it is beyond the scope of this chapter, as there are other grounds that are sufficient to defend the view that we have obligations to protect our own privacy.

For example, if one values one's own wellbeing (for utilitarian or other reasons), one ought to protect one's privacy because doing so helps to advance one's wellbeing (Allen 2013, 860). Protecting one's privacy will shield one from embarrassment, discomfort, security risks, discrimination, etc.

Crucially, there are also other-regarding reasons. One consideration is the nuisance of oversharing. Burdening non-intimate others with too much information about oneself may make them feel uncomfortable and overwhelmed.

More importantly, in a world avid for big data, every piece of information you give out about yourself helps create a map of profiles and types of personalities. The more you give out information about yourself to businesses, the more you are helping data brokers, not only construct a profile about you (which they will then

sell), but also construct profiles of people who are similar to you in relevant ways. Data brokers partly rely on correlations to make inferences and predictions about people in order to profile them; their aim is to have a profile on every Internet user around the world.

Anyone can buy those profiles. Customers of data brokers include financial institutions, insurance companies, cable and telecommunications corporations, political campaigns, retail stores, law enforcement agencies, possible employers, and sometimes, criminals. Profiles are not even expensive; bank account numbers are sold for 50 cents or a dollar (Dwoskin 2014) and have been sold to 'fraudsters' (Singer 2014). A full report on an individual can go for as low as 95 cents of a dollar (Angwin 2014, 7). Data brokers have also been known to sell lists of rape victims, AIDS patients, and more (Hicken 2013). These brokers can seriously cause harm by enabling would-be employers to discriminate people for reasons that should not be relevant to their employment, by undermining people's credit scores, by facilitating theft, by preventing people from hiding from potentially dangerous others (e.g., abusive spouses), etc. In 1999, for example, Liam Youens paid Docusearch, an online data broker that is still in business, to find information about a woman he was obsessed with. A few days later he found her at her workplace, killed her, and then killed himself (Angwin 2014, 7).[14]

[14] For more on data brokers, see the United States' Federal Trade Commission 2014 Report: "Data Brokers."
https://www.ftc.gov/system/files/documents/reports/data-brokers-call-transparency-accountability-report-federal-trade-commission-may-2014/140527databrokerreport.pdf

By protecting your own privacy, therefore, you are also protecting the privacy of others like you. Protecting others' privacy through your own is more direct in the case of family, as much is shared between nuclear family members (genes, location, socioeconomic status, etc.).

Showing restraint in exposing private information about oneself also helps to cultivate an environment where people can remain silent, if they want. A culture of exposure creates expectations that force people to share more than what they are comfortable expressing. In a worst-case scenario, cultures of exposure endanger dissenters and non-conformists. When people, for whatever motivation, constantly expose their normality—their having the 'right' sexual orientation, or belonging to the 'right' political party, or having the 'right' opinion—people who do not cater to the mainstream are put in a difficult position. Either they keep silent, and attract attention to themselves (which may mean endangering their lives), or they expose a false normality. Lying about these issues has at least two disadvantages: that one is being insincere about one's life, which may emotionally distance one from others, and the risk that one may be found out.

These considerations suggest that sometimes we should be sensitive about things that we would not be sensitive about if it were not for our concern for others' wellbeing. For example, suppose you live in a society ruled by a homophobic government. In this society there happen to be very few homosexuals. The information that you are heterosexual is not considered personal information in this context, as most (heterosexual) people are happy to share it with anyone in

order to avoid suspicion or wrath from their government, and far from being information that might hurt them, it is information that will keep them safe. In this context, however, heterosexuals should be sensitive about this information because they should be aware that by announcing their heterosexuality, they are failing to protect homosexuals. If many or all heterosexuals expose their heterosexuality, homosexuals will be singled out. It is as if heterosexuals were sharing personal information about homosexuals. At the very least, they are unwittingly and indirectly risking homosexuals' personal information, because the government will be able to recognise them by inference. If heterosexuals want to stand in solidarity with homosexuals, they will regard as sensitive all information about sexual orientation, thus making it personal information.

It is not always easy to protect one's privacy—for the sake of oneself and that of others—in the digital world. Most people are unaware of many of the data collection methods being used by governments and businesses. For example, businesses were having trouble tracking people across different devices (i.e., attributing a phone, a computer, and a purchase to the same person), so they have now developed a solution: ultrasonic cross-device tracking (Newman 2016). Inaudible sound beacons are being broadcasted through TV and radio commercials, as well as through music in shops, and are being picked up by customers' phones (without their knowledge), thus alerting businesses of their customers' location in order to better track their devices and purchases. These audio beacons function like sound cookies to triangulate one's devices and

purchasing habits through location. Few people are aware of this invasive practice.

Another little known technology through which privacy gets invaded are IMSI-catchers—fake cell phone towers that trick mobile phones into connecting to them. Once connected, IMSI-catchers can collect identification and location data, as well as permit eavesdropping on phone conversations, text messages, and web browsing (Schneier 2015a, 68). There is evidence that this equipment is being used by the police across London to spy on people, for example, at peaceful protests and near the UK parliament (Bryant 2016). Popular advice given online to protect one's privacy now includes leaving one's phone at home when going to a protest (Dunne 2017). Though they are mostly used by governments, IMSI-catchers could be used by anyone, as they are sold by private companies and can also be home built (Kolker 2016).

In addition to being ignorant about privacy risks, most people lack the technological abilities needed to use sophisticated privacy-protecting tools such as encryption. Moreover, successfully protecting one's privacy may come at a significant cost to oneself. It might mean not having a smartphone, for example. It may mean paying more money for certain services, which some people may not afford. There are also personal costs. People who live far away from their loved ones may not have the means to communicate with them in a secure way. They are forced to use the telephone, Skype, email, or Facebook—which are all privacy-insecure tools—to have their most intimate conversations with their close

ones. Morality can only ask so much of moral agents; people cannot be asked to refrain from telling their far away loved ones about their personal lives.

The many limits that individuals face in protecting their privacy suggest that institutions should pick up the slack. The discussion on the right to privacy already warranted the implementation of legal and institutional tools to protect our right to privacy. But the obligations mentioned here give the state and private organisations such as businesses further reasons to institutionalise mechanisms to respect privacy: they should aid us, or at the very least permit us, to fulfil our privacy-related obligations. If institutions do not offer us the means to communicate privately with close ones, for example, they are failing us twice: on the one hand, they are violating our right to privacy, and on the other, they are preventing us from fulfilling our privacy-related obligations.

IX. Conclusion

In this chapter I proposed a definition of privacy as remaining personally unaccessed. The value of privacy and the interests we have in keeping certain things private, which ground the right to privacy, were explored. I argued that the right to privacy should be understood as protecting a robustly demanding good. It was noted that perceptions of privacy should be taken into account in the moral exploration of privacy cases. Finally, I sketched some moral reasons to protect one's privacy and explored some of their implications. In the next chapter, I will

explore the relationship between security and privacy in the context of mass surveillance and terrorism.

CHAPTER FIVE

The Conflict between Security and Privacy in the Context of Terrorism and Mass Surveillance

When duties arising from two different rights are incompatible with one another, the rights in question can be said to be in conflict. Public discourse is flooded with claims about the incompatibility between the right to privacy and the right to security. According to popular belief, the more privacy individuals enjoy, the less the state is able to provide security, and vice versa. According to former NSA security consultant Ed Giorgio, '[p]rivacy and security are a zero-sum game' (cited by Wright 2008)—meaning that for every increase in one, there is a decrease in the other. In other words, the state seems to have incompatible duties: on the one hand, to respect its citizens' right to privacy by refraining from spying on them, and on the other hand, to guarantee its citizens' right to security, which, so the argument goes, cannot be done without spying on the general population.

In this chapter I focus on the supposed trade-off between privacy and security in the context of terrorist threats and mass surveillance. I understand mass surveillance to be the surveillance of the general population in an effort to collect as much data as possible from as many people as possible. The intelligence

community prefers to talk of bulk collection, denying that bulk collection amounts to mass surveillance, because most of the data collected is never analysed or accessed by any human (MacAskill 2015). If my account of the right to privacy is correct, however, the mere collection of data does violate the right to privacy because intelligence agencies like the NSA put themselves in a position where they can access that data at will. For this reason I side with organisations such as Privacy International when they claim that bulk collection of data does amount to mass surveillance (Falchetta 2016).

By the right to security I mean the right to be protected from harm (in particular, physical harm, but also financial and other kinds of harms). Throughout the chapter I will be talking of both the values of privacy and security (which their respective rights protect) and the rights themselves (as well as the duties they ground). On my account, the right to privacy protects more than just privacy—it also protects us from people securing a position from which to invade our privacy, and from attempts to invade our privacy. I believe the right to security also includes more than security—likewise, people should not achieve a position from which to attack us, nor should they attempt to attack us.

I will follow Waldron's (1989) framework for adjudicating between rights in conflict by first weighing security against privacy, thereby assessing the importance of the interests at stake (Section I). In Section II, I explore whether mass surveillance is a proportionate and necessary measure in response to terrorism and argue that it is not. Section III focuses on possible internal connections between

privacy and security that may suggest these rights are less in conflict than is usually thought. Section IV ends the chapter with some reflections on the implications for encryption. I will argue that encryption is not a significant obstacle to criminal investigations and therefore should be used widely.

I. The interests at stake—weighing security against privacy

Although the battle is still raging, it seems that in these times of fear, security is getting the upper hand in the choice between security and privacy. The last few years have witnessed many tragic terrorist attacks: from the attack on the Twin Towers in September 2001, to the Madrid train bombings in 2004, Breivik's attack in Norway in 2011, the Bataclan massacre in 2015 in Paris, the Baghdad bombings in 2016, and more. Terrorist attacks are traumatic events for populations to bear, and it is no surprise to find people concerned about security. In what follows I explore some of the reasons people appeal to when defending security over privacy.

Security: a collective good

It is common to think that, in the conflict between privacy and security, what is at stake is the right to privacy of the individual against the common good of the population. For instance, Sir Malcolm Rifkind, a former foreign secretary in the United Kingdom, has said that '[t]here is a balance to be found between our individual right to privacy and our collective right to security' (cited by Hopkins et

al. 2013). On this construal of the conflict between those two rights, someone who defends her privacy is portrayed as selfishly putting her personal interest before the common good.

This way of framing the debate is a misconception. As we have seen in Chapter Four, much like security, privacy is important both for individuals and for the common good, and is a crucial value for democratic societies. A society with no privacy would be an oppressive one (Solove 2011, 50). Privacy is important to protect freedom of speech and association, among other democratic values. A political activist seeking privacy to protect her activities is not doing it out of selfishness, but rather out of a desire to better society. We all have interests in living in a society where others and we enjoy both security and privacy, and framing the possible conflicts between these two rights in terms of the individual against the common good is misguided.

Security: more valuable than privacy?

It is also popular to believe that security is more morally valuable than privacy because, given a situation where a choice must be made between a loss of security and a loss of privacy, it is assumed most people would choose the latter. In the rest of this subsection, I focus on Kenneth E. Himma's arguments because they are a clear philosophical defence of this stance. This is not meant to be an ad hominem attack, however, and the view that security is more valuable than privacy is often found in the media, although perhaps articulated with less detail and clarity. Ron Iphofen (2016), for example, has written in defence of the superiority of security:

'What matters most in times of crisis? That people won't be able to find out who we are, where we are and when? Or that by accurate surveillance (…) disaster can be averted, lives saved and misery avoided?'

On Himma's view, '[o]ther things being equal, a security interest defeats a privacy interest that has the same level of moral importance to privacy that the first element has to security' (2016, 149). In other words: when the most important value protected by security conflicts with the most important value protected by privacy, security wins. If forced to choose between disclosing the most private piece of information about himself to save his life or save himself from grave injury, Himma says that he would 'gladly' and 'without hesitation' give up his privacy and that he 'would be surprised if there are any rational persons who would react differently' (2016, 152).

Whether Himma's choice is generalizable to 'any rational person' is not as evident as he assumes. For one thing, in some cases giving up private information may amount to giving up one's security: in homophobic societies, exposing oneself as a homosexual may seriously endanger one's life. (This is an example of the internal relations between privacy and security, which are explored in more detail in Section III.)

Furthermore, in some cases private information could wreck such havoc in one's life or that of others that some people might prefer to give up their lives or incur serious injury before exposing themselves. Consider again the case of

homosexuality in homophobic societies. The leak of sensitive photographs might not only damage the reputation of a person (and, perhaps, his professional life), but it might inflict serious shame and suffering on his family in certain traditional societies. One can imagine a person risking his life to prevent those pictures from being published, which shows that at least sometimes people choose privacy over security.

The most one can grant to Himma is that it may be the case that people *usually* tend to value their lives (or their physical integrity) more than they value their most private information. While that does say something in favour of considering security more valuable, it is not as black and white as Himma makes it seem. Furthermore and more importantly, we are rarely, if ever, faced with the choice proposed by Himma—the extreme choice. Both security and privacy come in degrees. More often than not, the choice is between having x degrees less of privacy for the benefit of y degrees more of security. That we may be inclined to choose security in the extreme choice does not imply that we would be willing to give up *any* degree of privacy for *any* degree of security: if the loss in privacy is enormous and the gain in physical security is minimal, then it is plausible to think that most people would prefer to preserve their privacy. Pointing out that security may weigh more heavily in the extreme choice of having to give up one's life is hardly helpful in making decisions in all those cases where the sacrifice in security is probabilistically small.

Himma also believes that privacy is less valuable than security because it is something we only value instrumentally, whereas we value security both instrumentally and intrinsically:

> If X is a right that protects something that is instrumentally valuable as a means to Y, something that is intrinsically valuable and protected by a right, Y is the more important value of the two from the standpoint of morality because the value of X derives from the value of Y in the following sense: but for the intrinsic value of Y, X would not be instrumentally important, and hence, would receive no moral protection. (Himma 2016, 154)

Whether privacy has intrinsic or merely instrumental value depends on how we conceptualise the distinction between instrumental and intrinsic value. For Himma, being intrinsically valuable means having value independently of one's instrumental worth, or having value 'as an end in itself.' By contrast, '[a]n entity has instrumental value if, and only if, it has value as a means to some other valuable end' (Himma 2016, 152). As mentioned in Chapter Four, Section II, another way of framing the distinction is in terms of consequences, as Raz does: 'Something is instrumentally valuable to the extent that it derives its value from the value of its consequences, or from the value of the consequences it is likely to have, or from the value of the consequences it can be used to produce' (1984, 205).

If we accept that the mark of intrinsic value is having value in itself, independently of consequences or of being a means to something else, Himma is wrong to think that privacy has no intrinsic value. While privacy may be mostly valued for instrumental reasons—because of the consequences it can have for physical

protection, for example—it also has intrinsic value. Consider the following two worlds: in the first world you enjoy privacy in your home. As a result, you enjoy peace of mind. In the second world, there is a Peeping Tom constantly gazing into your home, with no way of making him go away. Suppose that you know the Peeping Tom will not use any information he gleans about you against you, but you still feel very uncomfortable by his constant gaze. The only solution found is to take a pill offered by your psychiatrist that gives you peace of mind despite the intrusiveness of being watched all the time. It is intuitive to think that the first world is better than the second, even if the consequences are the same. There is some moral worth lost, apart from any consequences, when privacy is intruded upon. The analogy here would be the difference between obtaining pleasure from having a meaningful conversation with a friend versus obtaining pleasure from an experience machine that stimulates one's brain. Even if the felt pleasure was the same, it is intuitive to think that the former is more valuable than the latter.

We seem to value more the kinds of wellbeing that arise from certain cherished situations (i.e., the peace of mind enjoyed as a result of others respecting our privacy) than wellbeing arising from direct brain stimulation or deceit (i.e., experiencing peace of mind as a result of a pill that eases our concerns about privacy, or experiencing peace of mind thinking that we are enjoying privacy when in fact we are not because someone is surreptitiously watching us).[1]

[1] This reflection is in agreement with objective list theories of wellbeing, according to which wellbeing is constituted (at least partly) by other goods above and beyond pleasure or desire-satisfaction. Popular items include knowledge and friendship. For more on objective list theories of wellbeing, see (Crisp 2013).

In this sense, the intrinsic value of privacy is not dissimilar from that of security. It seems that the world is a better place if we are free from security risks, independently of whether we know about them (and therefore do not suffer their negative consequences, supposing these never come to be). In other words, it is better to have no one try to kill you, even if you never find out about it and the person does not succeed—just as it is better to have no one watching you, even if you never find out about it and no bad consequences ensue.

Tony Doyle (2009) has argued that there is nothing wrong with voyeurism that is never discovered, publicised, or exploited for other purposes. His perspective appears to be purely consequentialist, however. A purely consequentialist perspective does not seem to go well with the concept of intrinsic value—by definition, something valued above and beyond consequences.[2] If one accepts that there are things that are intrinsically valuable, that there is some wrong being committed by watching someone surreptitiously (or attempting to hurt her even if the attempt fails), and that the world is better without these wrongs, then it seems that both privacy and security have both intrinsic and instrumental value. Either way, even if one rejects this intuition and prefers a purely consequential approach, then it is still the case that privacy and security are on an equal footing. From a purely consequential perspective, both privacy and security are valuable in so far as they enable good consequences such as peace of mind (or wellbeing).

[2] From a classical consequentialist perspective, only wellbeing or happiness or pleasure are intrinsically valuable.

Even if privacy was merely instrumental and security was of intrinsic value (which, as we have seen, is not the case), it is not obvious that instrumental rights are always less important than other rights they protect or facilitate. First, there might be imbalances that betray the simplified formulation drawn by Himma. We might be dealing with different numbers on each side of the scale, and the severity of the loss on one side can be much graver than the other. As Himma himself acknowledges: 'It would not (…) be permissible to disclose the most private information of one thousand people to save one person from being bruised' (2016, 149).

Second, privacy is not *only* a means to physical security. It also achieves other benefits like financial and reputational gains, and, as has been mentioned already, relaxation and peace of mind, so that the weight of each benefit adds up to a point where privacy benefits can supersede the weight of certain security interests. As important as security may be, most of us are comfortable engaging in some physical risks for purposes such as convenience—consider how many people are willing to travel by car. Similarly, I suspect most people are happy to take risks that threaten their physical security in order to enjoy more privacy (with all its benefits).

Even if Himma's arguments are not convincing, the intuition that security always trumps privacy may still be felt strongly by many readers. Perhaps this is because security, while it may not be more *valuable* per se (it may not be always chosen over privacy and it may not be of more intrinsic value), seems more of a *basic* or

fundamental right than privacy. I now turn to Henry Shue's work on basic rights and explore security and privacy in this context.

Security: more basic than privacy?

Henry Shue has argued that the right to security—a right 'not to be subjected to murder, torture, mayhem, rape, or assault'—is a basic right, in that it is a necessary condition for the enjoyment of all other rights (1980, 13-34).[3] If one is at serious risk of getting gravely harmed by the police or of losing one's life, or indeed if one has already been hurt, one cannot be said to genuinely enjoy other rights such as that of freedom of association. If streets are so dangerous that one is afraid of going to school, one cannot be said to fully enjoy one's right to education. In order to enjoy a wide range of rights, the most severe impediments to exercising those rights have to be removed. Physical threats are barriers to the enjoyment of any right (Shue 1980, 21). Thus, guaranteeing physical security is a *sine qua non* for assuring other rights.

Physical threats are a powerful weapon against the defenceless, and an important function of basic rights is to 'prevent, or to eliminate, insofar as possible the degree of vulnerability that leaves people at the mercy of others' (Shue 1980, 30). According to Shue, given that security is a condition of possibility of other rights, if a choice must be made between security and some other right that is not basic

[3] Shue's conception of security is narrower than mine, as it only allows for physical security, but that has no bearing on my objections to his view. The same objections would apply to a broader conception of security.

(i.e., that is not necessary for the enjoyment of all other rights), security is to be given precedence (Shue 1980, 20).

At a first glance, privacy does not seem to be a basic right in Shue's sense because it is not necessary to enjoy *all* other rights. One can imagine enjoying certain rights, such as physical security, freedom of association, or education, without enjoying privacy (i.e., with the government, companies, or other citizens knowing who one associates with and what one studies). If the government in question is tyrannical, however, privacy may become more of a requisite to the enjoyment of at least *some* other rights. If one is afraid of suffering repercussions because of one's ideas or associations, then suddenly privacy is indispensable for enjoying a right to education, freedom of association and expression, etc.

It could be argued that one should not need to hide to exercise those rights, so that privacy is unnecessary. Even in the case where the government is not tyrannical and one is not in fear of repercussions, however, there is something to be said for the necessity of privacy to fully exercise one's right to education, freedom of association, freedom of thought, etc. Other people's gaze influences us in powerful ways. If we cannot prevent other people knowing what we read, whom we talk to, and what our opinions are, we are bound to change our behaviour. In a sense, then, privacy is indeed instrumental to the enjoyment of some other rights, even in democratic societies where people do not feel in danger.

Admittedly, however, privacy only seems to be a requisite to the *full* enjoyment of *some* political rights. Often, we do not fear the loss of privacy per se, but rather the negative consequences that may come with the loss of privacy: we fear for our physical security, the possibility of losing our jobs if our boss finds out something sensitive about us, or becoming stigmatised by our society, for example. As has been mentioned, however, there is intrinsic value in privacy. And even if privacy is thought to be instrumental, the importance of a right and whether it is basic or not need not be related to whether the substance of a right is intrinsically valuable or instrumental (Shue 1980, 20).

Notwithstanding these points, presumably Shue would still argue that privacy is not a basic right because it is not essential to the enjoyment of *every other* right. In particular, while physical security seems necessary to enjoy privacy, privacy is not necessary to fully enjoy physical security. One can imagine a situation where one is watched at all times by someone who poses no physical threat to one. Consider, for example, the case of an omniscient and benevolent god who would never harm anyone.

It is not entirely clear to me, however, that physical security is necessary to enjoy privacy. Consider someone who is thrown into an autonomous torture chamber (i.e., a torture chamber that does not need an agent to function). The victim has no physical security, but her privacy seems intact, as long as no one is watching and no one is prying on the victim's personal belongings or records. It may be that the victim is in such grave pain that she does not care that her privacy is intact;

perhaps using the verb *enjoy* to describe her status with regard to her right to privacy is inappropriate, as the victim cannot be said to be enjoying anything. Yet there is a sense in which she can be said to retain her right to privacy. And one could conjure up examples where the victim does care about it and has a strong preference for it to remain that way. Suppose the victim hides an important secret (say, a private piece of information that would shame her family if revealed) in a safety deposit box containing all of her personal information and no one knows about (you can suppose further that, even if the people who want to torture her know about it, they are adamant about respecting her privacy). Perhaps even in the darkest moments of being tortured, the victim can find relief in knowing that her privacy is safe (and her family safe from shame).

As the example above shows, I doubt that physical security is necessary to enjoy *every other* right. However, physical security is admittedly necessary to enjoy many more rights than the number of rights that need privacy to be fully enjoyed. In this sense, security is more basic, and therefore perhaps can be considered weightier. But the discussion above suggests that the moral distance between security and privacy in terms of how basic they are is not as large as one might think at first glance. In fact, neither is essential to the enjoyment of the other, and both often permit the enjoyment of many other rights.

Lessons in weighing security against privacy

So far we have seen that two out of three arguments to defend the superior value of security are successful, but much less so than their authors assume. The attempt to portray privacy as individual and selfish is misleading, as it fails to appreciate the ways in which privacy contributes to the common good. Privacy seems as intrinsically valuable as security, and while it may be true that people would usually choose to keep their lives over their most private information, we are rarely faced with such an all-or-nothing choice. Similarly, while security may be more of a basic right than privacy, both privacy and security enable the enjoyment of other rights, and it seems that we can enjoy privacy without security and security without privacy. In short, because the superiority of security is not as crushing as it might seem at first blush, we cannot defer to it in balancing privacy against security, because the latter does not always trump the former. More importantly, none of the three approaches seems very helpful to decide, in practice, how to adjudicate conflicts between security and privacy.

In the next section I will assess the trade-offs involved in security and privacy in the context of terrorism and mass surveillance.

In order to make an informed decision about whether to give up privacy, we should know how many people (and who) would have to give it up, how much of their privacy they would have to give up, and what will be gained in return for such a sacrifice.

Given that, thanks to Edward Snowden's revelations and the debate that has ensued since, we have more information on mass surveillance from the United States than from any other country, most of my empirical references focus on the American experience. The United States, however, is not alone in surveilling its citizens, and similar observations can be applied to other countries. Crucially, my arguments against mass surveillance are meant to criticise mass surveillance generally, and not the American system in particular—the empirical data offered simply illustrates the possible pitfalls of *any* system of mass surveillance. It just happens to be the case that the American system of mass surveillance is the closest we have come to a perfect system of mass surveillance (one that collects it all), and the one on which we have more evidence about how well or badly it has worked in fighting terrorism.

II. Proportionality and necessity

For mass surveillance to be justifiable, it must be proportionate and necessary. It has to be proportionate, in that its harms must not outweigh its benefits. And it also has to be necessary, in the sense that the same results cannot be achieved by less intrusive or harmful methods. These two conditions are independent: a measure could be proportionate (its benefits outweighing its harms), but unnecessary (if the goal pursued can be achieved in some other way), and necessary (it being the only way to achieve a goal) but disproportionate (its harms outweighing its benefits). Although independent, the concepts of proportionality

and necessity are both affected by effectiveness. For mass surveillance to be necessary, it must be effective, and the more effective it is, the more security gains it will have that will counter possible harms, and the more likely it is to be proportionate. In what follows I argue that mass surveillance is neither proportionate nor necessary.

Proportionality and surveillance

The moral concept of proportionality refers to a moral constraint on actions that cause harm. For an act that causes harm to be proportionate, it must be done in the pursuit of some valuable goal against which the harms are weighed (McMahan 2009, 19). If benefits outweigh harms, then the act is proportionate.

The idea of proportionality is not foreign to surveillance. The Investigatory Powers Bill in the United Kingdom includes proportionality among the considerations to be taken into account under general duties in relation to privacy. In 2014, 500 organisations and experts worldwide signed the International Principles on the Application of Human Rights to Communications Surveillance, which includes the principle of Proportionality (Schneier 2015a, 168). Proportionality has not been entirely absent from academic discussions about surveillance either (Brown and Korff 2009, Macnish 2015). However, proportionality has received much more attention, and has been developed conceptually much further in the field of Just War theory, from which I will be borrowing insights and applying them to the issue of mass surveillance.

Just War theorists typically recognise two requirements of proportionality: one pertaining to the decision of going to war (*ad bellum* proportionality) and one concerning the means used to fight in war (*in bello* proportionality). The first requirement maintains that waging war is impermissible if the bad effects outweigh the good, and the second holds that the same is true for individual acts of war (McMahan 2009, 19). Kevin Macnish (2015) has suggested there are parallels when it comes to surveillance. It can be said that *ad speculandum* proportionality governs the decision to employ surveillance, and *in speculando* proportionality concerns the methods of surveillance being used. While the distinction in Just War theory seems very relevant—the decision to go to war is a morally grave one, and *any* act of war is bound to create much harm—I find it less useful when it comes to surveillance. The decision to surveil (particularly as a response to terrorism) does not seem as morally grave, since the surveillance used could be relatively minimal (e.g., listening to only one conversation), and the degree and tools for surveillance used are much more morally relevant than the decision to surveil.

Whether surveillance is overall proportionate will depend on who gets surveilled, in what ways, for how long, etc. I will not be offering an answer to these queries. I will merely argue that *mass* surveillance is a disproportionate response in relation to the threat of terrorism.

Jeff McMahan also distinguishes between *narrow* proportionality, which refers to harm that is inflicted on those who are potentially liable, and *wide* proportionality, which refers to harm inflicted on those who are not liable to any harm (McMahan 2009, 21). In the context of terrorism, someone is potentially liable to be harmed by surveillance when they are implicated in a terrorist plot that could be stopped through surveillance. Someone is liable to be harmed if 'his own action has made it the case that to harm him (…) would not wrong him' (McMahan 2009, 11). In other words and for our purposes, because of their implication in terrorism, a liable person cannot justifiably complain of being harmed through surveillance.

I will be focusing on wide proportionality. I will not establish who, if anyone, is liable to be surveilled in the context of terrorism apart from terrorists (e.g., family members of terrorists). Given that mass surveillance affects *all* of the population, I will be concerned with the harms suffered by innocent people, who, I think it is safe to say, are a majority. Most people are not implicated in terrorism, under any reasonable definition of 'implication.'

I next intend to assess whether the harms of surveillance outweigh the good it produces. I will argue that they do by presenting an overview of the trade-offs we incur in giving up privacy for mass surveillance. Within the costs associated with losses in privacy, I will include the moral debts that arise from having the right to privacy unmet (what Waldron calls 'successive waves of duties').

After 9/11, public discourse in the United States often revolved around the idea of rethinking the balance between security and privacy (Waldron 2003). Unbeknownst to citizens, shortly after the attacks, a secret system of mass surveillance was implemented with the objective of increasing security. With the passing of the Patriot Act, six weeks after the terrorist attacks, the FBI was allowed to issue 'national security letters,' a form of subpoena that is not subject to judicial oversight and allows snooping into the private lives of people (phone records, bank accounts, web searches, and credit-card purchases) who might not even be considered suspects (Wright 2008).

In the past, implementing a system of mass surveillance was more complicated than it is today. The Stasi, for example, had to hire people to wiretap suspects, and pay for informers. They had one spy or informant for every 66 citizens (Koehler 1999, 9). Although complete information on all citizens is an intelligence agency's dream, the Stasi only managed to have files on roughly a quarter of the population (Angwin 2014, 40). Today, it has become much more feasible to hold information on most or all of the population. For one thing, a significant proportion of people volunteer private information in social networks. As filmmaker Laura Poitras put it, 'Facebook is a gift to intelligence agencies' (cited by Peterson 2014).

Furthermore, almost every citizen carries a potential spy in her pocket in the form of a mobile phone. Simply by accessing location data from mobile phones, one can know where someone works and lives, where they spend their free time and with whom, whether they go to church, whether they speed drive, and whom they sleep with (Schneier 2015a, 1). With this information, it is possible to predict where people will be 24 hours later, to within 20 meters (Schneier 2015a, 1). This information is not only used by totalitarian governments to track dissenters; democratic countries use it too. In 2010, Michigan police accessed mobile phones near an expected protest and without a warrant (Schneier 2015a, 2). At least two of the NSA's internal databases (HAPPYFOOT and FASCIA) contain location data of devices around the world (Schneier 2015a, 3). It is said that the NSA can even track mobile phones when they are turned off (Schneier 2015a, 3).

In the days of the Stasi, there was a significant chance of being the victim of surveillance, but one could still be lucky and not be spied on. Given that it has become cheaper and easier to eavesdrop on people, collect all possible data, and save it all indefinitely than to have to select relevant data (i.e., only from suspects), electronic mass surveillance means *every* citizen can be reasonably confident that sensitive data about her is being collected (Schneier 2015a, 23-28).

Accordingly, perhaps the most problematic aspects of the supposed trade-off between privacy and security in mass surveillance is that the right to privacy of *all* citizens is violated—even if they are neither criminals nor suspects—and the potential loss of privacy is significant—intelligence agencies today have the power

to access much more information than the Stasi ever did. In other words, bulk collection of data seems especially disproportionate because it surveils *everyone*, even when there is no suspicion of wrongdoing.

In Just War theory, the principle of discrimination forbids intending the deaths of noncombatants (though it allows the killing of noncombatants provided the harms are necessary, proportional, and merely foreseeable, not intended). Mass surveillance does not discriminate between liable and non-liable targets (people implicated in terrorism and people who have nothing to do with it). It could be argued that the real targets of mass surveillance are terrorists, and that the privacy harms imposed on other people are merely foreseen but not intended. However, snooping on the general population *is* intended and not merely foreseen—for there to be a mass surveillance system, *everyone* has to be surveilled. Documents provide good evidence for this. For example, in 2008, the Attorney General issued new guidelines allowing the FBI to carry out suspicionless investigations—investigations without 'any particular factual predication.' Similarly, in 2012, the National Counterterrorism Center got approval to copy government databases full of information on US citizens to look for suspicious behaviour (Angwin 2014, 27-28).

Traditionally, in most democratic countries, law enforcement had to give reasons for suspicion in order to apply for a warrant to surveil someone. Bulk surveillance allows law enforcement to go on fishing expeditions and *develop* grounds for suspicion (Schneier 2015a, 179). Yet the official view of most legal systems is still

that citizens must be treated as innocent until proven guilty (Angwin 2014, 18). Police officers typically have to have some reason for suspicion in order to stop someone on the street. With mass surveillance, however, every citizen is treated as a potential criminal—we are all suspects, even without reasons for suspicion.

The possibility of fishing expeditions by the police and intelligence agencies to find reasons for suspicion is all the more worrying given the scope of current mass surveillance. In the past, surveillance was about the present and the future—what someone was doing then and there and what she planned to do. When so much data gets collected throughout decades, we can now go into the past (Schneier 2015a, 35). In 2008, the NSA database called XKEYSCORE held metadata for a month; MARINA holds people's browsing history for a year; telephone metadata is held for five years. These storage limits are general guidelines. If an NSA analyst, however, finds something of concern, if you are a user of encryption, or if you have been unlucky enough to use certain keywords that are deemed suspicious, your data is saved indefinitely (Schneier 2015a, 36).

It is likely that, if you have enough information about someone, you can find him guilty of *something*. The risk is that the law can become arbitrary. It has been argued that, in the United States, everyone is probably a criminal, given how many behaviours are criminalised (Kozinski and Tseytlin 2009) and that people may be committing an average of three felonies a day (Silverglate 2011). If anyone can be considered a criminal, there is a risk that the government will persecute activists and dissidents who challenge authority.

Those are, roughly, the violations of the right to privacy incurred. There are also moral debts created by governments failing to respect people's right to privacy.

Successive waves of duties

When mass surveillance is implemented, the government fails to respect its citizens' right to privacy. Even if the spying were justified (which, I will argue, it is not), these failures must be addressed, as rights that have not been respected do not just disappear. As I indicated in Section II of Chapter Four, unmet duties create moral debts, or successive waves of duties (Waldron 1989, 512, 515).

The first duty that arises from violating the right to privacy of citizens would be to notify people of this breach. Governments should inform their citizens of their surveillance policies. Information does not have to be detailed enough to jeopardise criminal investigations, but the general public ought to have enough information to understand the scope and nature of government surveillance. Citizens of democratic countries have to be reasonably well-informed of what their governments do in order to fully participate in democracy.

If a government misleads citizens by keeping silent about mass surveillance or denying it, it is guilty of deception and manipulating people's perceptions of privacy (as we saw in Chapter Four), thereby failing to respect people's autonomy. Covert surveillance deceives victims about their world—it controls the victim's beliefs about whether she is being watched. These false beliefs affect the victim's

desires and her actions. She acts differently than she would if she knew she was being watched. Through deception, the government thwarts citizens' attempts to make rational choices for themselves. For people to be able to autonomously decide how to lead their lives, they must be reasonably well-informed about whether someone is watching them, among other things (Benn 1971).

Individuals who are special targets—whose information is not only being collected but also accessed—should be notified of this as well, except when a competent judicial authority finds that notice would harm a criminal investigation. In such cases, the targeted individual should be notified as soon as possible, as this is the only way targets can contest the information being held on them (perhaps the information is inaccurate) and the procedures used (perhaps there was wrongdoing involved in targeting that individual).

The second duty arising from violating people's right to privacy is the duty to minimise harms that might come about as a result of government spying. First, data should be kept safe. The many data breaches we have witnessed in the last few years suggest that governments and intelligence agencies do not have the necessary security to be trusted with our most personal information. In 2015, for example, hackers stole the private information (Social Security numbers, fingerprints, health and financial records, and more) of 21.5 million people from a database containing government background checks (Hirschfeld Davis 2015). The fact that Edward Snowden managed to steal more than a million documents from the NSA (Strohm and Wilber 2014), and his description of analysts having

unlimited access to any citizen's private data (Rusbridger and MacAskill 2014), are also causes for concern. Governments are responsible for creating a system of mass surveillance that stores sensitive data; as long as the data exists, there are risks to people, and governments are responsible for minimising those risks.

In addition to keeping data safe, governments infringing the right to privacy of citizens must guarantee good practices—rules and procedures that ensure abuses will not be committed. Some of the NSA's shortcomings in this respect have already been mentioned in Chapter Four.

Finally, people may have a right to some compensation for any wrong incurred in the government's infringement of their right to privacy. I find this proposal wanting, as it is unclear what kind of compensation would be appropriate for the wrongs of mass surveillance, but perhaps some kinds of compensation would be satisfactory. At the very least, reparative justice may demand an apology and an explanation.

At the moment, governments such as those of the United States and the United Kingdom have not complied with these derivative duties: they do not inform people of privacy breaches, evidence suggests they do not keep data safe enough, and they do not offer any kind of compensation, apology or explanation for violations of the right to privacy. I do not intend to argue that, if governments were to meet these derivative duties, then it would be morally acceptable for them to engage in mass surveillance. If mass surveillance were disproportionate, as I will

argue it is, it would still be morally impermissible. Not meeting these moral debts, however, makes the wrongs of violations of the right to privacy all the more egregious.

Even with these losses and unattended moral debts in mind, it could be argued that, although violations of the right to privacy are significant for every citizen, the trade-off is still worth it, as the increase in security also applies to every citizen. The people asked to make a sacrifice are the same people who will benefit from it, thus making the sacrifice more palatable. An increase in security is compensation enough. Although each person has their right to privacy violated, each person can be confident that she has a better chance of avoiding being the victim of a terrorist attack. In order to determine whether harms outweigh benefits in the trade-off between security and privacy, possible gains in security must also be explored, which in turn calls for an assessment of the damage caused by terrorism. I now turn to this issue.

The threat of terrorism

Even though the threat of terrorism may loom large in many people's minds, when looked at from a numbers perspective, the threat seems significantly less salient. Presumably, the main objective of preventing terrorist attacks is to save lives.[4] If you are an American, however, you are equally or less likely to die from a

[4] Some people may think that the main objective of mass surveillance is rather psychological in nature: it is to make citizens *feel* safe. There is no reason to think, however, that mass surveillance is more effective than other methods (e.g., putting

terrorist attack than from lightning, allergy to peanuts, or being crushed by your own furniture (Solove 2011, 43, Zenko 2012, Plumer 2013). More worryingly, Americans are nine times more likely to be killed by a police officer than by a terrorist (Schneier 2015a, 135), and, because of loose gun control legislation, toddlers kill more people than terrorists do (West 2016).

The eight deadliest terrorist attacks in the history of the United States add up to fewer than four thousand deaths. In contrast, sixty thousand people die from flu and pneumonia every year (Solove 2011, 43). Other threats like car accidents and health problems arising from obesity are much more dangerous for any given citizen than a possible terrorist attack, yet many more resources are spent in trying to prevent the latter.

Given these statistics, it does not seem that we are in a particularly violent moment in history where extreme measures might be justified. Terrorist attacks and attempts have become considerably less frequent since the 1970s both in the United States (Plumer 2013) and in Western Europe (Merelli 2015, Stanley 2016). People living in these parts of the world are safer than ever. Moreover, even without mass surveillance, technology has provided police and intelligence agencies with more information than ever about possible suspects (as we will see later on in this chapter), so that they are better placed than they have ever been to investigate suspects and capture criminals.

---

the number of deaths into perspective, showcasing the effectiveness of targeted surveillance, etc.) for assuring citizens. Furthermore, this is not the justification given by governments for mass surveillance. Governments in Western countries typically claim that mass surveillance is necessary to ensure the security of citizens.

Given the limited damage that terrorist attacks have perpetrated in the last decades, and the huge scope of privacy invasions effected through mass surveillance, it seems that bulk surveillance is a disproportionate response to the threat of terrorism. It could still be argued, however, that mass surveillance is justified if it can prevent a massive terrorist attack—one that could end the lives of millions of people. If one imagines a big enough attack, perhaps the proportionality balance can weigh in favour of mass surveillance, after all. Even if we grant this view for the sake of argument, for mass surveillance to be justified it would still need to prove itself effective and indispensable (i.e., necessary) to fulfil its preventative role. In what follows I argue that mass surveillance is ineffective in fighting terrorism and that targeted surveillance works better.

The ineffectiveness of mass surveillance

The ideal mass surveillance system is one that collects and analyses all possible data. According to an NSA slide presented at a 2011 meeting of five nations' intelligence agencies, their goal is to 'collect it all,' 'process it all,' "know it all,' and 'exploit it all'. The UK's GCHQ also mentions collecting it all in a 2010 document (Greenwald 2014, 97).

But is it essential to have an all-watching system to prevent terrorist attacks? When giving up civil liberties such as privacy for the sake of security, we must be sure that the sacrifice will in fact have the intended consequence (Waldron 2003, 208). The state should show that mass surveillance is *necessary* to combat terrorism.

Necessity is closely related to effectiveness. If mass surveillance is not effective, or not as effective as other measures, it can hardly be necessary. In turn, the more effective mass surveillance is, the more benefits in security it will have that will outweigh privacy harms, and thus the more chances it has of being proportional.

One would think that with increased information, more terrorist attacks would be prevented. Two weeks after Edward Snowden's revelations, United States President Barack Obama claimed that 'at least 50 threats had been averted' thanks to information provided by mass surveillance programmes (Isikoff 2013). In 2013, Keith Alexander echoed the claim by affirming that 54 attacks had been thwarted. No evidence has been provided to support these claims, however (Elliott and Meyer 2013). In fact, the evidence that has been gathered and publicised, and the accumulated experience since 9/11, tell against the effectiveness of mass surveillance in the prevention of terrorism.

The reports that have looked at mass surveillance in the United States have all suggested that those methods were not necessary to prevent attacks. The President's Review Group on Intelligence and Communications Technologies observed bulk surveillance on telephone records 'was not essential to preventing attacks,' as the relevant information 'could readily have been obtained in a timely manner using conventional' orders, and that 'there has been no instance in which NSA could say with confidence that the outcome would have been different without' the metadata programme (Clarke et al. 2013, 104, 120). One of the members of the panel, Geoffrey Stone, a University of Chicago law professor,

admitted that '[t]he results were very thin.' The panel investigated whether the mass collection of telephone call records had actually stopped *any* terrorist attack: 'We found none,' said Stone (cited by Isikoff 2013).

Similarly, US Judge Richard Leon, who ruled the bulk collection programme to be unconstitutional, said that the government was unable to point out 'a single instance in which analysis of the NSA's bulk (…) metadata collection actually stopped an imminent attack, or otherwise aided the Government in achieving any objective that was time-sensitive in nature' (cited by Isikoff 2013).

A 2014 report by New America Foundation that analysed 225 cases of charged terrorists inside the United States since 9/11 corroborates the findings cited. It concludes that the bulk collection of phone records 'has had no discernible impact on preventing acts of terrorism,' and that traditional law enforcement and investigative tools are usually able to provide the necessary evidence to initiate the case against terrorists (Bergen et al. 2014, 1).

In 2015, The New York Times won a Freedom of Information Act lawsuit asking for the government to declassify a 2009 report by five agencies' inspectors general about the Stellarwind programme, which was composed of warrantless bulk phone and email collection activities. In 2004, the FBI analysed the tips that were gathered from Stellarwind to see how many had made a 'significant contribution' to identifying terrorists, deporting suspects, or developing a relationship with an informant about terrorists. Only 1.2% of the tips from 2001 to 2004 were useful

(Savage 2015a). When the FBI looked at the data from 2004 to 2006, they found that no tips had made a significant contribution (Savage 2015a).

The tips sent from the NSA to the FBI were too many, and too much of a waste of time (Savage 2015b, 162-223). When the NSA complained to the FBI that they were not seeing results from the information they were passing on, an FBI official responded that the feedback he was getting from field officers was: 'You're sending us garbage' (Report on the President's Surveillance Program 2009, 637).

We are persuaded to accept surveillance on the pretence that it will keep us safe, but the lack of evidence as to the effectiveness of systems of surveillance for preventing terrorism is striking. In the weighing of harms and gains towards establishing proportionality, so far it seems that, while losses on the privacy side in mass surveillance are considerable, the gain in security vis-à-vis terrorism is nil or negligible at best.

Still, one might think that the fact that mass surveillance has not had impressive results during the past 15 years in the United States is no reason to give up on it. Michael Morell, former director of the CIA and a member of President Obama's Review Group on Intelligence and Communications, is of this opinion. Protesting against the press's focus on the lack of effectiveness of surveillance programmes, he insisted on the value of these for national security. Morell believes that the bulk collection of phone metadata has the potential to prevent the next 9/11, and argues that '[i]t needs to be successful only once to be invaluable' (Morell 2013).

One might think that American intelligence officers in the past 15 years have simply been unlucky in the lack of effectiveness of their mass surveillance, and that it is only a matter of time before their efforts bear fruit.

There is reason to believe, however, that mass surveillance can never be effective in preventing terrorism. In what follows I suggest that mass surveillance, in virtue of its design, is not the best method to prevent terrorist attacks—targeted surveillance is both more effective and carries less risk of harm.

The superiority of targeted surveillance

A common misunderstanding in public debate is to think that, because someone is opposed to mass surveillance, she is opposed to any kind of surveillance. There are important differences, however, between targeted (also called traditional) surveillance and mass surveillance that are worth remarking on here.

Mass surveillance is the indiscriminate surveillance of the whole population. With targeted surveillance, the police or intelligence agencies must first get a tip. The tip usually comes from the community or from family members, but it can also come from informants, or relevant information can turn up from other criminal investigations. Once there is a genuine reason for suspicion, investigative authorities present the evidence to a judge, who must decide whether there is enough evidence for suspicion to order a warrant for surveillance to take place. This system works surprisingly well. Because police officers will have to go back to

the same judges for warrants in future cases, they are careful to build trust with them and ask for warrants only when it is quite likely they will find something. In fact, police find at least some of the evidence they had expected in more than 80% of cases (Solove 2011, 130).

Experience since 9/11 suggests that targeted surveillance is more effective than mass surveillance in preventing terrorism, which supports the conclusion that bulk surveillance is neither necessary nor proportionate for achieving our counterterrorist goals. The major technical problem with mass surveillance seems to be that which characterises it: the extent of the collection of data. As security expert Bruce Schneier puts it, 'when you are watching everything, you're not seeing anything' (2015a, 137). The sheer quantity of information adds irrelevant data about innocent people and obscures what would be significant tips from targeted surveillance. Mass surveillance adds hay to the haystack and makes it all the more difficult to find the needle.

Schneier has argued that bulk collection and data mining are inappropriate tools for finding terrorists for three reasons (2015a, 136-140). First, error rates are unacceptably high. When you use data mining to target people for something relatively innocuous such as fashion advertisement, mistakes can be tolerated more easily, as getting advertisements on clothes we do not want to buy is usually not too problematic. But when data mining is used to find terrorists, the lives and freedom of potentially innocent people are at stake, and our tolerance for mistakes should be low. Consider airports. More than two million people fly daily around

the world. If a data mining program to identify terrorists is mistaken even 1% of the time, then more than twenty thousand passengers will be flagged unjustly every day (Solove 2011, 188). Being flagged could lead one to be on a no-fly list (with no possibility to appeal the decision), which can cause significant and unjustified harm to an innocent individual. That terrorist attacks are infrequent compounds the problem.

Terrorist attacks are very rare in comparison to, say, fraud or other criminal activities that happen on a daily basis and create patterns that can be studied more accurately. The rarity of terrorism makes it the case that prediction systems will suffer from false positives—cases where harmless people are labelled as dangerous (Schneier 2015a, 137). The cost of false positives is very high. Alerts require further investigation to determine whether the target is in fact dangerous. It costs time, money, resources, and attention from intelligence officers. In the years following 9/11, the NSA passed thousands of leads per month to the FBI, all of which turned out to be false positives (Bergman et al. 2006). False positives are also very costly to victims, who are unfairly treated as suspects or criminals.

The second reason for the inappropriateness of data mining techniques for investigating terrorism is that terrorist attacks are unique. There was no way of predicting that pressure-cooker bombs would be used in the Boston Marathon attack (Schneier 2015a, 138), or that someone might put a bomb in his shoe, or that people would try liquid explosives, or that someone might kill people with a cargo truck in Nice. As a result of the unique methods of past attacks, the police

targets what would otherwise be catalogued as normal: a woman got a visit from the police after researching pressure cookers online (Gabbatt 2013), we all take our shoes off at the airport, and we cannot carry liquids in our hand luggage. As more attacks or attempts happen, and the list of unique methods grows, our liberties will be thwarted in unpredictable and ineffective ways: 'Each rare individual who carries out a terrorist attack will have a disproportionate impact on the criteria used to decide who's a likely terrorist, leading to ineffective detection strategies' (Schneier 2015a, 138).

Jeff Jonas, an IBM research scientist, and Jim Harper, the director of information policy at the Cato Institute agree with Schneier. In a paper they wrote together they conclude that 'terrorism does not occur with enough frequency to enable the creation of valid predictive models' (Jonas and Harper 2006, 8). A report by the National Academy of Sciences also concludes that, despite their success in the commercial sector, 'highly automated tools and techniques cannot be easily applied to the (…) problem of detecting and preempting a terrorist attack, and success in doing so may not be possible at all' (National Academy of Sciences 2008, 2).

Finally, the third problem is that terrorists will be trying to avoid detection, making it harder for them to be caught in the very broad nets cast out by intelligence agencies. It is reasonable to think that wrongdoers will be particularly well informed concerning the kinds of patterns that intelligence agencies are

looking for (e.g., people buying pressure cookers), and will avoid those and, again, choose novel and unique methods.

It seems, then, that mass surveillance is not as effective in preventing terrorism as other, less intrusive measures, which supports the claim that it is neither necessary nor proportional. In the balancing of privacy and security in mass surveillance, so far it appears that the violations of the right to privacy are steep and the gains in security are shallow. In a case of conflict, then, it is far from clear that security should be given precedence. But perhaps privacy and security are not as incompatible as it seems at first glance. In what follows I explore the internal relations between privacy and security.

III. Internal connections

As we saw in Chapter Four, in order to resolve conflicts between rights, sometimes the most appropriate strategy is to establish the relative importance of the interests at stake and then maximise what we deem of most importance (Waldron 1989, 518-519). Often, however, attention to internal connections may be more suitable for the resolution of conflicts. It is important, then, to recognise when certain interests undermine the right they purport to claim (as in the case of Nazi speeches that undermine the form of society that enables free speech, as indicated in Chapter Four). It has already been pointed out how privacy can be important for physical security. In what follows I explore in more depth some of the internal

connections between privacy and security, particularly when it comes to the digital age.

During the last few years the amount of information online about each of us has been dramatically increasing. For both professional and personal reasons, more of us spend more time on the Internet, using more websites and applications for our everyday activities, and creating data about ourselves. We have intimate conversations over messaging apps, buy groceries online, access our bank through our smartphones, and carry out a significant amount of our work-related tasks in front of a screen. The more time we spend online, and the more information we give up to websites, the more we become exposed to cyberharms, which include identity theft, blackmail, scams, data leaks, and more. While the NSA might be willing to acknowledge that their activities undermine people's privacy, they are adamant that this is done to protect security. However, there are ways in which intelligence agencies' practices are endangering the public's security.

First, the NSA makes the Internet insecure by stockpiling 'zero-day' vulnerabilities. Cyberharms are enabled through vulnerabilities—mistakes or cracks in the design of websites that allow intrusion. Both criminals and intelligence agencies exploit vulnerabilities in order to hack their way into computers and steal passwords, eavesdrop, and so on. Exploiting vulnerabilities is useful for collecting and accessing data. When someone discovers a vulnerability, she can alert the relevant institution so that it can be patched and its details published so that others can learn from the experience. She can also keep it to

herself to exploit it now or in the future, or she can sell it to some other hacker who wants to take advantage of it. Unpublished vulnerabilities are called 'zero-day' vulnerabilities.

At the moment, the NSA has a policy of stockpiling zero-day vulnerabilities (Schneier 2015a, 146). Possessing these enables the NSA to snoop at will whenever they want to, taking advantage of cracks in systems in order to access sensitive information. This policy puts everyone at risk. Instead of vulnerabilities being patched as soon as possible to protect people, they are purposefully left exposed, allowing hackers to discover them independently and attack websites. As a result of these attacks, personal information such as credit card numbers, addresses, and passwords are stolen and sold to the highest bidder.

Second, the NSA makes us all unsafe by inserting 'backdoors' (or deliberately created vulnerabilities) into commercial IT software and hardware (Schneier 2015a, 147). The problem with vulnerabilities is that they not only allow access to the government—anyone who finds them can exploit them as well. In 2010, for example, Chinese hackers exploited a vulnerability Google had put into Gmail to comply with the US government (Schneier 2015a, 148). 'As technologists,' says Schneier, 'we can't build an access system that only works for people of a certain citizenship, or with a particular morality, or only in the presence of a specified legal document. If the FBI can eavesdrop on your text messages' so can criminals and terrorists (2016, 2).

Third, intelligence agencies like the NSA and the FBI try to influence policies that undermine encryption. James Comey, the director of the FBI, has been a staunch critic of commercial encryption, arguing that it creates a safe haven for criminals (Ackerman 2014).

Fourth, the NSA routinely hacks the Internet. It breaks into computers and equipment, it pretends to be different websites to redirect users to its dummy sites, it carries out cyberattacks, and more (Schneier 2015a, 148-149). Sometimes, innocent users are harmed as collateral damage, for example, by being accidentally infected with a virus meant for someone else, or by suffering blackouts.

Finally, there is a tight relationship between information and power. Whether knowledge is to be thought of as being an instrument of power, or as a form of control over others in itself (Foucault 1977, 27-29), it is clear that the more knowledge someone has about us, the easier it is for them to hold power over us. Diminishing civil liberties such as privacy enhances the power of the state, and decreases the protection individuals may have against state incursions, even if it may enhance security against terrorism (which, as we have seen, is doubtful, at best) (Waldron 2003, 195). This is particularly worrisome given that there is some evidence that intelligence agencies such as the NSA at times seem especially concerned with the security of the state, making sure they are not the target of damaging public debates, and so on. Noam Chomsky has suggested this is the case in a public conversation had with Edward Snowden and Glenn Greenwald. He

claims that in the documents leaked by Snowden, 'almost nothing is concerned with the security of the population' (Chomsky 2016). The documents show, rather, the NSA's concern with economic spying (Schneier 2015a, 73), diplomatic advantages (Ball 2013), and keeping what the government is doing secret from the domestic population. It seems, then, that intelligence agencies like the NSA may have power as part of their main motivation to have a surveillance state. With the wrong government in place, an effective system of mass surveillance could be lethal for political dissenters: it could provide the foundation of a near undefeatable dictatorship.

These are only some of the ways in which a lack of privacy undermines security. There are other examples in the world beyond the Internet, some of which have been mentioned in past chapters. What all of them show is that sometimes there is a relation between privacy and security such that the less one has of one, the less one has of the other: Ed Giorgio was wrong to think that security and privacy are a zero-sum game.

Further evidence is the fact that not all measures that enhance security undermine privacy. For example, one of the many security changes that were implemented after September 11 was to lock the cockpit doors on airplanes (Solove 2011, 34). At other times, there are two ways of improving security, and one of them will be less invasive of privacy. For instance, there is some evidence suggesting that streetlights are as good at deterring crime as surveillance cameras (Welsh and Farrington 2004), and the former can be more privacy-friendly than the latter (it

depends on who obtains access to cameras, whether footage is recorded, for how long it gets stored, etc.).

So far we have seen that, while security may weigh more heavily than privacy in some extreme circumstances and enable the enjoyment of more rights, these observations on their own do not help us make decisions when it comes to real choices we have to make in the face of security threats such as terrorism. In the trade-off between privacy and security, violations of the right to privacy seem significant and gains in security seem scant, making mass surveillance disproportionate and unnecessary. I have further established that privacy and security are not always in conflict. Given their internal connections, less privacy often leads to less security. This means that we do not always have to choose between them. In the next section, I will explore some of the implications that my arguments have for encryption, since it is a controversial topic that has gained much attention these days, and since it may have a significant effect on governments' abilities to engage in mass surveillance.

VI. Implications for encryption

The recent rise in encryption practices has given intelligence agencies another excuse to ask for more powers of mass surveillance. The FBI and other intelligence agencies have been lobbying hard against encryption, arguing that it gives criminals a chance to hide their wrongdoings, and that their investigative capabilities are 'going dark' (Olsen, Schneier, and Zittrain 2016, 6-7).

Evidence suggests, however, that encryption does not hamper criminal investigations in significant ways. In most cases, police and intelligence agencies can break encrypted communications and devices. An example is the San Bernardino case in 2016: Apple refused to create a backdoor into the iPhone of terrorist Syed Farook, so the FBI ended up paying an alleged $1.3 million to a hacker to gain access into the phone (Yadron 2016). In 2014, there were only four cases in the United States where the police did not succeed in breaking encryption in wiretaps (Greenberg 2015). If those few cases had been important enough, the police would have probably found the necessary funds and capable hackers to break into the relevant devices. As Schneier puts it, '[i]f a sufficiently skilled, funded, and motivated attacker wants in to your computer, they're in. If they're not, it's because you're not high enough on their priority list to bother with' (Schneier 2016, 2).

If police decide not to pay the price to break into a device, it is good to keep in mind that authorities have a large array of investigative tools at their disposal: among other useful avenues, in the United States and most of Europe, police can access data stored in the cloud and metadata (in some countries they do not even need a warrant).

Most metadata cannot be encrypted and contains very useful information. Metadata is data about data. It is information that computers need to operate and it is the by-product of such operations. It includes information such as the

operating system that created the data, the time and date of creation, the author of the data, and the location where the data was created. In a telephone call, for example, the contents of the call are considered data, and the metadata include the phone numbers of both the caller and the person receiving the call, date and time, the duration of the call, and location data.

Metadata may seem uninteresting, but one can glean quite a lot from it: where someone has been, what they did, who they talked to and for how long, what they read, what they purchased. Researchers from Stanford University examined the phone metadata of volunteers. From it, they could deduce highly sensitive information. For example, they could tell that someone had multiple sclerosis (he had called local neurology groups, a special pharmacy, a rare-condition management service, and a hotline for a pharmaceutical dedicated to treat multiple sclerosis); that another participant was a marijuana grower (he had called a home improvement store, locksmiths, a hydroponics dealer, and a head shop); and that someone had had an abortion (she had had a long call with a family member, then a series of calls to a Planned Parenthood clinic) (Mayer and Mutchler 2014).

Metadata is here to stay. Many systems need it to operate, and most of it cannot be encrypted, so it seems that law enforcers would have plenty of opportunities to carry out criminal investigations (and even mass surveillance) even if encryption were to be adopted as a general practice (Olsen, Schneier, and Zittrain 2016, 3,

9). This is particularly the case with the coming of the Internet of Things, when all devices turn 'smart' and begin gathering information about users.

Furthermore, there are other available sources of information for law enforcement. Surveillance can easily be carried out in cloud backups, which are typically not encrypted (if they are, the business owning them usually has the keys, so that they can be turned over to law enforcement agencies). Moreover, encryption does not protect one's devices when in use. If a device is encrypted, it cannot be used, and if it is in use, it is unencrypted. That is how the FBI gained access to Ross Ulbricht's data. Ulbricht ran the biggest online illegal drug marketplace: Silk Road. While he was running the site in a library, with his laptop open and logged into the site, the FBI created a diversion to distract him and physically snatched his computer away (Bertrand 2015).

In short, encryption cannot deter skilled snoops from accessing private data. Even with widespread encryption practices, law enforcement, companies, and criminals can all collect significant amounts of sensitive metadata. However, encryption does protect content (not metadata): it makes the job of surveilling content difficult enough to force listeners to target (Schneier 2016, 2). In other words, it is an effective measure against the mass surveillance of content. Police and intelligence agencies will still be able to surveil the general population through metadata (unless there are laws and effective oversight to impede it), but if they want to access content, they will need to choose the content of a few, rather than surveil everyone.

As we have seen, however, targeted surveillance seems to be more effective than mass surveillance, and more than enough surveillance for the purposes of criminal investigations. It is more expensive, but still affordable, as it is cheaper than in the past. Hiring personnel to follow someone covertly in foot or by car costs around $175,000 per month. In the past, this was the only possibility available. Placing a tracker in the suspect's car or using a device to access the person's phone at a distance would cost the police about $70,000 per month. If the police manage to hide a GPS receiver in the target's car, then the price is only $150 per month. The cheapest option, however, is getting information from the suspect's mobile phone provider—only $30 per month (Schneier 2015a, 25). Only in special cases when law enforcement has to break into encrypted devices is it more expensive, which forces authorities to choose their targets in terms of priorities.

VII. Conclusion

While security may weigh more heavily than privacy in some extreme circumstances and in virtue of enabling the enjoyment of more rights, careful analysis of the trade-offs involved and the internal connections between privacy and security suggest that it is unreasonable to hold security as a supreme good that always trumps privacy.

In the case of the use of mass surveillance to prevent terrorism and the conflict between privacy and security, it seems that violations of the right to privacy are so

great and gains in security are so thin that the trade-off is disproportionate. Surveillance should be permitted only to pursue the gravest threats to the security of citizens, and the state must prove that such surveillance is necessary and proportionate for achieving its legitimate objective. I have argued that mass surveillance does not fulfil those requirements. Instead, targeted surveillance, which infringes the right to privacy of much fewer people and only when there is reasonable suspicion, seems to be a much better method of preventing terrorism.

Furthermore, the internal connections between privacy and security suggest that these values are not always in conflict, and that often a decrease in privacy entails a decrease in security. If Shue is right that a vital function of rights is to minimise the 'degree of vulnerability that leaves people at the mercy of others' (1980, 30), then it appears that we should be much more preoccupied with protecting privacy than we currently are.

The arguments explored in this chapter suggest that mass surveillance (of both data and metadata) should be stopped. In democratic societies, the ideal would be for mass surveillance to be replaced through appropriate legislation with targeted surveillance. However, there are at least three reasons in favour of developing technological solutions in addition to regulation. First, legislation is slow, and cutting edge technology always seems to be several steps ahead of it. Second, it is hard for effective laws to be passed to curb surveillance when governments benefit so much from it. Third, even proper legislation might not be able to control hackers and criminals, intelligence agencies that work underground, and foreign

businesses and governments over which they hold no jurisdiction. To compensate for the shortcomings of legislation, technology such as encryption can mitigate the harms of mass surveillance. Although encryption cannot stop the mass surveillance of metadata, it can minimise the scope of mass surveillance when it comes to content, and it is not the impediment it has been claimed to be for authorities in their effort to catch criminals.

CHAPTER SIX

Radical Transparency: A Thought Experiment in Surveillance, Sousveillance, and Coveillance

Governmental mass surveillance, collection of personal data by companies, and data breaches of many kinds are currently among the main menaces to privacy. As I write this chapter, not a day goes by in which concerns about privacy do not surface in the press. Just this month, it has been reported that a data breach at the Office of Personnel Management in the United States has compromised personal information of about 21.5 million people (Sciutto 2015), and a freedom of information website in the United Kingdom said it has recorded 154 accidental data leaks made by public bodies since 2009 (Ramesh 2015).

Distress about privacy threats is felt among ordinary people, old and young, on both sides of the Atlantic. A 2014 Pew Research Center report affirms that in a survey done in the United States, 91% of adults agreed or strongly agreed 'that consumers have lost control over how personal information is collected and used by companies' (Madden et al. 2014, 3). In a subsequent report, Pew found that 88% of Americans surveyed said 'that it is important that they not have someone watch or listen to them without permission,' and 65% 'believe there are not adequate limits on the telephone and internet data that the government collects'

(Madden et al. 2015, 18, 31). In the United Kingdom, research done by the think tank *Demos* suggests half of young people are either extremely or very concerned about online privacy—in 2014, they were more worried about privacy than about environmental issues, immigration, tax avoidance, or the EU (Birdwell, Cadywould, and Reynolds 2014).

For people concerned about current trends, there are two general strategies to take in order to change the privacy climate being experienced and resist surveillance. The first and most common approach is to try to safeguard privacy by building metaphorical or literal walls. If you want to avoid your neighbour from spying on you in your garden, build a higher wall. If you want to avoid others reading your emails, encrypt them. If you want to pass by unnoticed, hide. According to this approach, if we want to feel safe, we must build laws that act as barriers to limit government surveillance. Privacy advocates of this sort typically believe that through walls one may protect not only one's privacy and that of others, but also other positive values, such as freedom, that are believed to be supported or enhanced through privacy.

The second strategy one may want to take to uphold freedom, peace of mind, and maybe even privacy in a world of surveillance is the opposite to the one just mentioned. It is based on the idea of taking down walls, banishing the shadows by letting the light shine in, building society out of glass instead of bricks.

For some, a crucial premise is that we cannot escape being watched (Brin 1998, 8-9). If surveillance—usually, organisations observing people—is here to stay, and fighting it is futile or not productive enough to make the battle worth it, then the best we can hope for is to watch the watchers, demand full disclosure from all parties. If policemen are going to watch us, we might as well turn our smartphone cameras on and watch them back. Such inverse or bottom-up surveillance has been dubbed *sousveillance*—watching from below (Mann, Nolan, and Wellman 2003). If we add *coveillance* (equals watching each other) to surveillance and sousveillance (Mann, Nolan, and Wellman 2003), the result is a transparent society that, if taken to its extreme, permits any individual or organisation to 'see' what any other individual or organisation is doing. 'The answer to the whole privacy question is more knowledge,' says Kevin Kelly, the founder of Wired magazine: 'More knowledge about who's watching you. More knowledge about the information that flows between us—particularly the meta information about who knows what and where it's going' (cited by Quittner 1997).

In this chapter I explore the proposal of extending transparency as much as possible—letting the sun shine into every corner of life. I consider what a radically transparent society would be like, and analyse the ways in which such a society would be better or worse than alternatives. In Section I, I present the supposed virtues of transparency. In Section II, I explore the proposal of radical transparency from an optimistic point of view. In Section III, I present some limits to the virtues of transparency and possible problems a radically transparent society would face. From the perspective of privacy, a radically transparent society would

be better than alternatives if either: a) it is easier to protect privacy in such a society, or b) it has so many other advantages that the benefits outweigh the sacrifice in privacy. I will argue that neither condition obtains. Section IV concludes the chapter by detailing some guidelines as to when transparency can be inappropriate.

I. The supposed virtues of transparency

Transparency—'the principle of enabling the public to gain information about the operations and structures of a given entity' (Etzioni 2010, 389)—is a highly regarded value in contemporary Western societies. A positive connotation accompanies organisations and people who can be described as transparent. Presumably, the more transparent an entity is, the less wrongdoing it is able to hide, and the more trustworthy it is. In the words of Onora O'Neill:

> There is quite a large measure of consensus about the way that transparency is supposed to work. It is supposed to discipline institutions and their office-holders by making information about their performance more public. Publicity is taken to deter corruption and poor performance, and to secure a basis for ensuring better and more trustworthy performance. (O'Neill 2006, 76)

Strong defenders of transparency believe that it can replace government regulations (e.g., Crovitz 2009). The assumption is that, if transparency allows people to see the inside workings of entities, those entities will automatically avoid transgressions, thereby making government controls superfluous.

It appears that transparency counts among its advantages two main virtues. First, it seems to facilitate holding people accountable. Second, it enables people to make informed decisions.

Accountability

If all actions in an organisation, for example, can be traced back to individuals, responsibility can be accurately attributed: people who do things right (both morally and professionally) can be praised, people who make mistakes can be taught how to avoid them in the future and may be reprehended (where appropriate), and people who commit wrongdoings can be blamed and face legal consequences (where they are called for). Accountability through transparency, then, assists justice in providing accurate information to give each person her dues. (Institutions may also be held accountable qua institutions, of course.) Accountability can also prevent wrongdoing and make people feel safe.

On the assumption that adequate and effective transparency practices are put in place, and all concerned parties are aware of such practices, it is reasonable to think most potential wrongdoers will refrain from misbehaving. Jeremy Bentham thought it 'an indisputable truth' that 'the more strictly we are watched, the better we behave' (2001, 277). Given his conviction, it should be no surprise that Bentham was the designer of the famous panopticon—an institutional edifice that permits a single watchman to observe everyone inside the building without people being able to tell if they are being watched at any particular moment (though they

are aware that they *might* be watched at any time). Although the paradigmatic panopticon is a prison, the design can be used for other kinds of buildings where surveillance is important, such as hospitals, schools, and work-houses (Bentham 1995, 34).

Two centuries later, in his *Discipline and Punish* (1977), Michel Foucault would use the panopticon as a metaphor for modern 'disciplinary' societies and their tendency to surveil and normalise behaviour throughout all contexts. With present day technologies of recording and data storage, the panopticon can be enhanced: instead of making people believe they *might* be observed at any one time, contemporary panopticons can give people the certainty they *are* being watched at *all* times and places.

Presumably, most people who cheat and steal do it because they think they can get away with it. Stealing is not an attractive option if one will get caught before one can even enjoy what one stole. If people had a high degree of confidence (or even certainty) that transparency practices are such that they would get caught quickly were they to misbehave, they would probably avoid temptations. The exact mechanisms through which people (and organisations) would refrain from wrongdoing are presumably related to an absence of gain in wrongdoing, as well as the desire to avoid punishment and shame, and maintain a good reputation. Good reputations are good for business, and for escaping the bad consequences resultant from people feeling moral outrage against one.

Confidence in accountability through transparency can also make people feel safer. The head of a company would feel safe in the knowledge that her employees are not stealing from her. In turn, employees can feel safe in the knowledge that management is not cheating them. A consumer can feel safe in the knowledge that she is eating what the label on the food product says it contains. If the police were subject to transparency, citizens could feel safe in the knowledge that their rights will be respected. And so on. In other words, it seems that through transparency, people can be more trusting of organisations and other individuals.

Perhaps this sense of safety, in addition to being related to a true decrease in wrongdoing and deceit, is also psychologically related to what Steve Mann (2013), the inventor of wearable computing, considers a 'basic human need' of 'being able to see better.' Darkness disturbs us. Even when we are fairly certain we are safe, experiencing darkness—not being able to see for ourselves what is in front of us—can be distressing. Being kept in the dark, literally and metaphorically, is not something most people appreciate in most circumstances.

Facilitating informed decisions

Good markets and democracies both depend on people's informed decisions. Each decision to buy one product instead of another is a 'vote' to keep a product in the market, and a company in business. For consumers to make good choices about what to buy, information is crucial. Food labels that include information about sugar contents, for example, enable people to make better decisions for their

health and for supporting the kind of products they want to see more of in supermarkets. Information about corporate social and ecological responsibility included in labels enables consumers to 'vote' on the kind of products and companies they approve of.

Similarly, with more transparency, people will be able to make more informed decisions about who to hire based on having extensive background information on all applicants for a position. Having hired someone, clients can be confident in their choice to maintain that professional relationship as a result of having all the relevant information about their employee's performance.

Democracy also benefits from having well-informed citizens. Standard mechanisms of democratic participation and oversight include elections, public opinion, and public deliberation, all of which depend on citizens having access to information (Sagar 2007). The more information people have on the government and on politicians, the better they will be able to judge them and vote according to their preferences.

Given the purported virtues of accountability and facilitating informed decisions, it is not hard to understand why transparency has few enemies, and why anyone arguing against it can be perceived as suspicious. It is easy to find instances where it is relatively clear we want at least *some* degree of transparency. Examples include audited financial statements by corporations, campaign contribution disclosure, food labels (ingredients, calories, GMOs, etc.), warning labels on hazardous

materials, and privacy policies. Questions remain, however, about whether transparency is *always* a good thing. Are there contexts where transparency is undesirable? And even in settings where it is desirable, is there such a thing as too much transparency? In the following section I consider a completely transparent society in order to assess these questions.

II. Radical transparency: surveillance + sousveillance + coveillance

The inevitability of transparency

Although in theory one could favour a radically transparent society even if one thought surveillance could be avoided and privacy protected, most transparency advocates seem to think that there is no way of stopping surveillance (one-sided transparency). David Brin, for example, thinks that 'it is already far too late to prevent the invasion of cameras and databases. The djinn cannot be crammed back into its bottle. No matter how many laws are passed, it will prove impossible to legislate away the new surveillance tools and databases. They are here to stay' (1998, 8-9).[1, 2] The choice, then, is not about having cameras or not having them, but rather about who will ultimately control the cameras.

[1] Daniel Dennett and Deb Roy (2015) agree that our future will be a transparent one, and that organisations will either adapt or perish. They caution, however, that transparency is a 'mixed blessing.'

[2] Stuart Armstrong (2013a) also thinks that privacy laws cannot be trusted, that protecting one's privacy in the usual ways is futile, and that focusing on extending transparency to be able to know how your information is being used is 'a much more achievable goal.'

According to Brin, if we were to pass privacy laws against cameras and against people knowing certain things about other people, the laws would be broken by the mighty (presumably, both governments and corporations), and an underground market of knowledge would be created (1998, 198). Citizens have little reason to trust intelligence agencies to follow the law. Governments have a tendency to collect as much information as possible to have as much control over the population as they can. If citizens disagree with these policies, intelligence agencies can always hide what they do, just as the NSA did. Similarly, personal information might be too valuable for companies to give up. As security expert Bruce Schneier points out, '[s]urveillance is the business model of the internet' (cited by Gillmor 2014). Brin's conclusion is that privacy laws are useless, and that we should pursue radical transparency as a second best bet.

Transparent paradise

David Brin begins his book, *The Transparent Society* (1998), by asking his readers to imagine two cities. Both cities are technologically sophisticated, and street crime has nearly disappeared from both towns. Tiny cameras observe all pedestrians in both cities. In city number one, the camera footage gets sent to police or government headquarters, where security officers or 'agents of some mysterious bureau' process images, identify citizens, and act in consequence. In contrast, in city number two, the camera footage can be accessed by *anyone* through their

smartphone.[3] 'Here a late-evening stroller checks to make sure no one lurks beyond the corner she is about to turn. Over there a tardy young man dials to see if his dinner date still waits for him by a city fountain. A block away, an anxious parent scans the area to find which way her child wandered off' (Brin 1998, 4).

In the second city, cameras are banned from some indoor places such as homes, but not from police headquarters. Cameras accessible to all citizens would be placed in the police camera control room to make sure authorities are on the look out for crimes, and only crimes. Throughout the book, Brin argues that city number two is the better world.

Although ubiquitous and accessible-to-all cameras would be of utmost importance in a transparent society, they are also just a symbol for all kinds of data collection. For radical transparency to be achieved, society would need to go beyond visual information. Given that many people develop most of their activities on their computers, every keystroke, website visit, and changes to documents would need to be recorded, and the information would also need to be accessible to any citizen and organisation.

The condition that ordinary citizens may monitor organisations (sousveillance) is meant to undercut the primacy of organisational surveillance. Mann, Nolan, and Wellman argue that '[s]ousveillance disrupts the power relationship of surveillance when it restores a traditional balance that the institutionalization of Bentham's

[3] Brin's book was written before the age of smartphones, so he calls the devices used to access cameras 'wristwatch televisions.'

Panopticon itself disrupted' (2003, 347). It can neutralise surveillance by levelling the informational playing field and restoring equality between surveillers and surveillees.

A real-life example can illustrate the power of sousveillance to benefit equality and justice. In 2013, a police informant planted some crack in a small shop in Scotia, New York (Engel 2013). The police then recorded themselves acting as if they had found the crack and proceeded to arrest the owner of the shop, Donald Andrews Jr. Luckily, Andrews had footage from his own security cameras that showed how the police informant had previously planted the drugs. Andrews was subsequently released and cleared of all charges. Had Andrews not had the same power as the police to record the events of that day, he would probably be in prison. The danger with authorising only one party of a transaction to record what happens is that, in the event that the recording will not suit that party's needs or preferences, the recording may be lost, ignored, or shown only partially.

For this reason, Mann criticises businesses and institutions that have a policy of prohibiting people from entering their premises with cameras. He gives an analogy for the importance of symmetry in surveillance in what he calls 'The Veillance Contract' (Mann 2013). Suppose, he asks, that you sign a contract with a corporation. The corporation does not allow you to take a copy of the contract with you (no photography or note taking is allowed either). A few years later, the corporation sues you for contract infringement. When you get to read the contract in court, the content is different from that which you remember signing, but you

have no way of proving that the contract was changed. This thought experiment emphasises the injustice behind organisations banning sousveillance but engaging in surveillance.

Mann believes that if A makes a recording of a transaction it has with B and forbids B from doing the same, then A should not be able to use that recording as evidence in court. For Mann, people should be able to carry around cameras everywhere in order to protect themselves, just as shopkeepers are entitled to have cameras to protect their products.

According to Stuart Armstrong (2013b), 'if all goes well,' apart from equalising power differentials and eliminating most crimes, radical transparency would bring about:

A decrease in police force and powers

Because crime would go down, a large police force would not be needed anymore. Furthermore, there would be no need for police to enjoy some of the special powers they are currently entitled to—searching people, arresting people on suspicion, interrogating them, lying to them during interrogation (in the United States), etc. Today's police forces are given those powers because they enable them to acquire information. The police do not know whether a suspect is armed; therefore, they must search him. The police are not sure whether a suspect is guilty of a crime; therefore, they must interrogate him. In a radically transparent society, however, these powers would be dispensable. For the most part, guilt or

innocence would easily be determined through recordings of various types from the comfort of an office chair.

Fewer laws

According to Armstrong (2013b), most countries have an excess of laws that is tolerated only because excess laws are enforced either rarely or selectively. If everyone were subject to enforcement in virtue of radical transparency, 'there would have to be a mass legal repeal. (…) When it becomes glaringly obvious that most people simply can't follow all the rules they're supposed to, these rules will have to be reformed.' As a result of fewer laws, we could expect an increase in personal freedom.

Reduced armies and warfare

Armstrong (2013b) argues that, in the past, international attempts to reduce armaments have been thwarted by a lack of reliable verification. If radical transparency were instantiated internationally, countries could be sure that other governments will not be deceiving them. If intelligence agencies were confident that other countries were not preparing for an attack or building new weapons, countries could shrink their militaries. With reduced armies, reduced warfare would likely follow.

Prevention of global catastrophic risks such as pandemics

One of the most worrisome risks to humans is that of a global pandemic. In the case of pandemics caused by human action, as a terrorist attack (engineered pandemics), or as a result of an accident in a research lab, radical transparency

could help prevent a global pandemic by spotting danger early: catching people who are attempting to engineer a pandemic and alerting the appropriate authorities of inadequate safety standards or dangerous practices in laboratories. In the case of naturally brought about pandemics, radical transparency could catch symptoms at an early stage (or perhaps detect diseases even before symptoms kick in, if we have sophisticated enough technology to monitor people's health). Thanks to surveillance and coveillance, we could quickly track all the people with whom infected individuals have come in contact with in order to conduct early quarantines and limit the spread of the disease (Armstrong 2013b).

An increase in academic freedom

In a radically transparent society governments could prevent the terrorist use of nuclear and other futuristic weapons. Therefore, it would be unnecessary to prohibit dangerous research projects or the flow of information regarding how to build weapons. As soon as a weapon began to be built in an unauthorised manner, the police could confiscate it. With less to fear on the part of governments, there would be an increase in academic freedom (Armstrong 2013b).

Research advances

Along with expanded academic freedom, research would benefit from the massive amount of big data that would be produced through radical transparency. Current methods of research usually take a sample of the population as the object of their study and take it to be representative of the wider population. Often, however, the sample in question is not representative for reasons of gender imbalance, race, culture, etc. With big data, every single individual can be

included in the data gathered. We can expect the accuracy of big data to dramatically increase the quality of research on humans. As Armstrong puts it, '[t]he [democratised] panopticon would be a research nirvana' (2013b).

An increase in convenience

With radical transparency in place, locks, passwords and other security measures would be unnecessary. Gone would be the days of carrying keys around, remembering dozens of passwords, and wasting time in security queues at airports. The surveillance system could be used for accurate identification of people (and detection of prohibited items) if it tracked all individuals at all times. Armstrong hypothesises that perhaps there would be no need for credit cards or cash registers. The camera system could bill people by identifying who they are and what they took from the shop (much like what Amazon is already trying out with its checkout-free shop in Seattle). Drivers who crashed into parked cars would not need to leave a note, since tracking would enable insurance companies to take care of the matter automatically (Armstrong 2013b).

An increase in truth-telling

In a transparent society, there would be no opportunities for lying and hypocrisy. Trying to give a false image of oneself would be futile and counterproductive (Armstrong 2013b).

III. The limits of transparency—problems in paradise

I imagine many people would find the transparent society so described attractive. The description offered so far, however, has only covered the possible advantages of a transparent society. Disadvantages must likewise be taken into account. In what follows I question the power of transparency to bring about its two main virtues—accountability and facilitating informed decisions. I then go on to look at the darker risks of trying to create a radically transparent society.

Accountability through transparency

For transparency to yield true accountability, a great deal is needed above and beyond transparency that often gets neglected by transparency enthusiasts.

Transparency, by itself, merely requires *disclosure* or *dissemination.* Data, however, can be disclosed in ways that make it useless. As Onora O'Neill has argued, certain epistemic and ethical standards must be met before transparency can yield trustworthiness (on the part of the individuals or organisations providing information), and trust (on the part of the general public) (2006). These standards, argues O'Neill, must not be taken for granted. They are far from trivial, and upholding them will be much more of a challenge than mere transparency: [4]

First, data has to be made available in ways that make it intelligible to its intended

[4] O'Neill does not number what the relevant ethical and epistemic requirements are, so this list should be read as my interpretation of her views and arguments.

audience. Effective communication requires that data not be cluttered with irrelevance, that the most important information be highlighted. O'Neill believes that the main purpose of transparency should be to communicate information. She worries about corporate practices that turn what should be a mode of disclosure into disclaimers (O'Neill 2006, 88). Businesses often use transparency to reduce risk by transferring liability to their clients. Risk is shifted from those who provide products and services, to those who purchase them. A good example of this practice are privacy policies that, apart from giving full powers to companies, are often thousands of words long and are written in legalistic jargon that most people would not understand even if they had the time to read them.

Transparency may thus lead us to a model of society where it is considered morally acceptable for people to suffer bad consequences as long as they were previously informed of the risks; in other words, a model that values autonomy more than the protection of people (Etzioni 2010, 403).

Second, information must reach relevant audiences. The reason why requiring big companies to have audited accounts works, is that company members are certain that someone with the relevant expertise *will* look at those accounts and scrutinise their performance. It is not enough for information to be accessible if there is a good chance nobody will actually access it. Independent entities must be appointed to go through the data produced by transparency (O'Neill 2006, 82-83).

Third, information must be true: it must be accurate (in that it does not contain

false truth claims) and honest (in that people are not making claims they believe to be false). Elsewhere, O'Neill worries that people who know they will be scrutinised will 'massage the truth' through hypocrisies, half-truths, self-censorship and deception (2002, 73).

Similarly, Rahul Sagar (2007) argues that

> the fundamental flaw in proposals to increase transparency [in the government] is structural in nature because their success is destined to rely upon the faithfulness of officials, which is ironic since the point of the whole exercise is to prove rather than assume their good faith.

I agree with O'Neill that the first two standards (making data informative and making sure it reaches an appropriate audience) are crucial challenges for an expansion in transparency to be successful. The third standard, however, seems less worrisome in a *radically* transparent society. If transparency relies on surveillance, sousveillance, and coveillance, it seems safe to assume that it will not need to rely heavily (or perhaps it will not have to rely at all) on individuals' reports or their truth telling. If *every* step, conversation, and keystroke is recorded and reported, we will likely need to trust whoever or whatever program processes that information into a report that highlights what is relevant, but not the individuals who are the protagonists of said reports.

The role of mediators, however, could be problematic—especially if the information being disclosed is too technical for the general public to understand. To be able to trust mediators, we must be fairly certain that they are getting *all* the

relevant information, that they are competent in processing that information, and that they are honest (i.e., that they are not distorting information to fit any interests). Transparency advocates would surely want to remind us that mediators too would be subject to scrutiny. But we risk a problem of infinite regress if we have to appoint someone to oversee the mediators, someone else to oversee the overseers, etc. I will come back to this issue because it is even more relevant for the challenge of cognitive overload.

Beyond O'Neill's requirements, a fourth and much more important requisite to make transparency bring about accountability is having a proper rule of law. Exposing people who commit crimes is not enough to bring them to justice. In countries led by corrupt oligarchies, uncovering wrongdoing or suspicious behaviour often leads nowhere. In 2014, for example, the Mexican journalist Carmen Aristegui revealed that the Mexican President Enrique Peña Nieto's £4.4m mansion (his residency), is officially owned by a company associated with some of the most lucrative contracts the government has offered (Holman 2014). As journalist Jorge Ramos (2015) pointed out in a gala where he was recognised by *Time* as one of the most influential people in 2015, 'in any country with a little bit of rule of law, the president would have been forced to resign. Guess what happened? The president of Mexico didn't resign, and the journalists who denounced the corruption were fired.'

In conclusion, transparency by itself provides an incomplete basis to bring about accountability. As a minimum, ethical and epistemic standards are also needed, as

well as a proper rule of law. These requisites can complete what can amount to a full system of checks and balances. Other parts of the system will have to include already common practices such as professional certification, complaints procedures, and meritocratic appointment procedures (O'Neill 2006, 87).

Facilitating informed decisions through transparency

For transparency to enable people to make informed decisions, O'Neill's requirements must also be met: information must be intelligible, true, and steps must be taken to ensure that it reaches its intended audience. Even then, while it can be said that transparency so described makes it possible for informed choices to take place, there is no guarantee that people will make better decisions.

First, there is the problem of information overload. In a society with full transparency, we would be so flooded with information that it is doubtful we would have the energy and time to profit much from all of it. After a long overview of studies, Susanna Kim Ripken concludes that '[e]vidence suggests that when people are given too much information in a limited time, the information overload can result in confusion, cognitive strain, and poorer decision-making' (2006, 160).

It has been estimated that if an average American internet user were to read all of the privacy policies of the sites she visits in a year, it would take her around 244 hours (worth about $3,534), or a whole month if she dedicated 8 hours a day only

to that task (McDonald and Cranor 2008, 563-564). Most people do not even have the luxury of having that many days for vacations, let alone for reading privacy policies. And the estimate only takes into account privacy policies. Even if these were made dramatically shorter and simpler to read, if we add transparency information about salaries, investments, sources of raw material, corporate responsibility (including information about ecological footprint), etc., of every business we come in contact with on a daily basis, and similar data about governments, it is clear it would be impossible for an ordinary individual to process such quantities of information. Nowadays, people are lucky if they have enough time to get through all of their emails on a given day.

Information overload could further inequalities. Those with enough resources will be able to hire people to process information that is relevant to them. Even if the result of this processing of information were made available for others to see, what is relevant to person A is not necessarily as relevant to person B, which means that the person who commissions sophisticated analyses of information will still get more out of transparency than an ordinary citizen.

Second, even if we could somehow manage to narrow down information to digestible doses for individuals, it is unclear how much people would profit from it. Amitai Etzioni doubts that transparency will have a significant effect on decisions (2010, 402). Even if there is no controversy in the raw data, processing information involves interpreting data, and there is almost always more than one way to do it. Statistics is a controversial discipline, and different ways of

calculating probabilities will have very different outcomes. In addition, even if we had perfect information, cognitive limits and biases such as overconfidence, loss aversion, and risk aversion cloud our judgments (Kahneman 2011)—not to mention forces such as targeted advertising, which may worsen our cognitive biases. At the very least, more evidence is needed to show that people would in fact take significant advantage of the data available to make better and more informed decisions, and that those benefits outweigh the costs of information overload.

One may be tempted to think that both these problems would disappear with the help of intermediaries, experts, and technologies that process information. Instead of having to review the information themselves, people can always rely on certification icons such as TRUSTe—a provider of privacy certifications for online businesses. Here again, however, as Etzioni points out, 'issues we face in dealing with the absorbability and veracity of what might be called first-order information we also face when dealing with intermediaries, that is, with second-order, processed (rather than directly accessed) information' (2010, 400).

The public does not seem to have the cognitive skills or the time to determine the quality of intermediaries or interpret their rankings and reports. In the United States, hospital rankings have been issued based on data compiled by the Centers for Medicare and Medicaid. On the basis of such rankings, people may understandably wish to avoid hospitals that score high in mortality rates. But, however unintuitive we might find it, these are likely the best hospitals because

they attract patients with severe illnesses (Etzioni 2010, 401). Similarly, it is unclear how people should deal with conflicting rankings (e.g., different university rankings). Even if the rankings were to be transparent about how they come up with their lists, and even if experts were to issue opinion pieces commenting on the merits and demerits of each ranking, in the likely case of there being controversies, it is unclear how ordinary people should make a decision.

Certification seals are also not without problems. In 2014, TRUSTe was fined $200,000 by the United States Federal Trade Commission after being accused of misleading users to believe it was doing proper auditing of the companies to which it gave its stamp of approval (Fox-Brewster 2014). Another example is the 'USDA Organic' label, which has been criticised for lowering their standards and allowing substances such as synthetic additives to count as organic as a result of being pressured by lobbyists (Kindy and Layton 2009).

In short, the public would need intermediaries to process the vast amount of information that a radically transparent society would produce; yet they will also have reason to doubt those intermediaries and they would have trouble interpreting the processed information. If people cannot judge matters for themselves, or follow, review, and challenge information disclosed, transparency will have failed to bring about informed decisions and trust in institutions (O'Neill 2006).

What this subsection shows is that, once again, transparency by itself is not

enough. More proposals need to be offered regarding how to regulate intermediaries, experts and technologies that process information. To my knowledge, radical transparency advocates have yet to delve into such details.

Even though I have argued that transparency has important limits in bringing about accountability and facilitating informed decisions, from the point of view of someone worried about privacy, a radically transparent society might still be a better option than alternatives if it can better protect privacy or if it has so many other advantages that they outweigh sacrifices in privacy.

Privacy in a radically transparent society

Even though Brin argues in favour of a radically transparent society, he still wants to preserve privacy. 'Indeed,' he writes, 'without some privacy, we could scarcely function as humans' (Brin 1998, 14). He thinks of transparency as a tool for protecting privacy (299), and believes that privacy is a product of liberty. Therefore, if we manage to have a transparent society that respects liberty, Brin assumes privacy will come for free (3, 78-79, 201).[5] He admits, however, that it will not be privacy as many might currently dream of it. Rather, it will be 'a little'

[5] An in-depth exploration of the relationship between privacy and liberty is beyond the scope of this chapter. It is worth noting, however, that even if it is empirically true that privacy is often the result (and maybe also partly the cause) of a free society, freedom and privacy do not *necessarily* go together. Someone can have privacy without having freedom. A person may hold another captive in the latter's home without ever violating her privacy: the captor wears earplugs and never looks through the window; she simply does not let her prisoner go out. Conversely, someone can have freedom without having privacy. Consider the possibility of an omniscient god who never interferes with people.

privacy—what he calls 'bedroom privacy' (13). For Brin, essential privacy is that of home (he calls homes 'sacrosanct'), hearth, and the intimacy one shares with loved ones (9, 26).

It is unclear, however, how bedroom privacy could be preserved in a transparent society. Brin mentions how, in his transparent city, cameras would be 'banned from some indoor places' (4)—presumably, from homes. Perhaps this comment can be explained by the fact that Brin wrote his book roughly at a time when webcams were only starting to become popular. Now that cameras are included by default in a number of electronic and personal objects such as laptops, smartphones, TVs, and game consoles, the idea of the home being a camera-free 'sanctuary' seems less feasible.

Even in 1998, however, Brin seems anything but naïve. He repeatedly mentions the possibility of cameras the size of insects ('gnat cameras,' he calls them) and the difficulty of controlling them (e.g., 271, 286). He hopes that the development of gnat cameras may prompt the invention of 'antignats'—devices that patrol one's home to detect unwanted cameras (286). A captured gnat camera could be hacked to learn its point of origin, says Brin, and make it transparent.

Detecting when one's own smartphone or laptop camera has been compromised, however, may prove even more challenging. If one adds all the microphones that can listen to us to the equation, the difficulty of noticing when one is being snooped on becomes apparent. Samsung smart TVs capture people's

conversations and send them to third parties, Facebook can turn on your smartphone's microphone, and iPhone's 'Hey Siri' and Android's 'OK Google' features listen all the time (Schneier 2015b). Given that our data gets sent out to third parties, determining who is listening to us is a huge challenge. Even if there were laws requiring disclosure of listening third parties, it is unclear how we could be sure that some third parties will not be kept hidden. When we bring our smartphones into our bedrooms, the private and public spheres become so intertwined that drawing a privacy line at the door of one's home seems like an impossible thing to do in a transparent society.

To be fair, to carry out one's most private activities in peace, one can always turn off one's laptop and smartphone, put them in the freezer (as Edward Snowden asked reporters to do when they first interviewed him (Murphy 2013)), and build a bunker bedroom with no windows. This option is as much of an alternative in the present as it is in a future radically transparent society. It seems to me, however, that in a transparent society, with people in general having much more access to knowledge about oneself, one would need to take many more precautions to enjoy a sliver of privacy with peace of mind. Home would not feel as homey if one needed a bunker bedroom with antignats flying around in order to enjoy privacy. Even then, for the ordinary person, it would be very hard to hide who one brings into one's bunker bedroom—and that means having much less privacy than most people would probably be comfortable with.

Detecting intrusions through technology and exposing the perpetrators is not

Brin's only answer to the challenge of how to maintain privacy in a transparent world. Courtesy, thinks Brin, will be crucial. To illustrate the importance of cultural norms, Brin offers as an analogy a restaurant (1998, 14-15). There are usually two ways of having a private conversation with someone: either one finds a quiet and solitary place, or one finds a public place. In public places such as restaurants, it is rude to stare at others, and most of us frequent these places with the relative certainty that others will not violate this norm. There is no need to make laws requiring people to wear blinkers or requiring restaurants to erect paper screens between tables. In fact, Brin believes it would be worse to require the latter, because it would encourage people to poke little holes on the screens in order to stare with impunity. In screen-free restaurants, mutual civility and the desire to avoid embarrassment by being caught in the act of staring are sufficient to keep people from minding others' business. According to Brin, the same incentives could work in all contexts of a transparent society:

> [A] culture of openness will sustain some privacy, if that is what free citizens want, and if Peeping Toms have reason to fear getting caught. Courtesy may return as an important moderating force, for the simple reason that it will make life among the cameras more bearable—and because those who don't practice it will be found out, losing their neighbor's good will.
> I believe this balance of technology with common sense may result in a world where we are observed only about 80 percent of the time (…)
> Above all, citizens will be much too busy to spend time peering at one another. They'll have better things to do. (Brin 1998, 301)

I find Brin's optimism misplaced. Even if the world he is describing is in the realm of logical possibility, more needs to be said, and more evidence is needed, to make his world seem like the likeliest of outcomes once a transparent society is built.

Other, much darker options seem likelier. I have two main objections.

First, even if it were true that we would remain unwatched for 20 per cent of our daily lives, we might not get to choose which 20 per cent of our lives we can keep to ourselves. If our smartphones can be accessed by anyone, just like the cameras on the street, believing that one will probably be unwatched and unheard for about 4 or 5 hours a day might be of little comfort—particularly if those hours are the hours we are asleep, or doing something where we would not mind others watching. An important aspect of privacy is being able to *choose* it at certain times and with certain people. Knowing that one could be watched at *any* time has almost the same psychological consequences as being watched *all* the time.

Second, we should not underestimate people's curiosity and thirst for gossip and voyeurism. Gossip is as old as language, and there is no evidence people will get bored of it any time soon. Consider how successful reality shows and gossip TV shows are. On the contrary, our tendency to gossip can intensify considerably with increased access to people's intimate lives. Similarly, from the armchair, one might think (quite reasonably) that people would get bored from watching porn if they had unlimited access to it. As Internet trends show, however, appealing to people's most basic instincts does not seem to grow old. Imagine if people had facilitated access to people's homes (through knowledge of their address, access to their smartphones, etc.). I find it hard to believe an important number of people would not take advantage of that facilitated access for the purposes of voyeurism and gossip.

Brin is not always as optimistic as in the quotation above. At one point he admits that '[b]usybodies will gossip, but,' he adds, at least 'you'll know *their* secrets' (1998, 334). Knowing other people's secrets in exchange for giving up one's own may not be of much comfort for the shiest among us, as well as those who are less prone to prying into other people's lives. Having an exhibitionist watch one in the intimacy of one's home in exchange for being able to watch him back does not sound like a desirable prospect for most people (I hope).

A transparent society does not seem to fare well with respect to privacy. One could argue, however, that it is unfair of me to measure the desirability of a transparent society in terms of privacy because a transparent society is one in which privacy is no longer needed. Privacy helps us maintain security and other goods that would be available in a transparent society.

In Chapter Four, we saw that privacy is not the only way to achieve some of the things it typically provides. In a completely transparent society where getting away with theft, kidnap, or murder would be impossible, privacy with respect to our financial assets and location would no longer be needed to keep our bodies and money safe.

When it comes to embarrassment, not feeling judged by others, and being free from the anxiety of being looked at, however, privacy does not have any substitute. One might think that in a transparent world people will simply get used

to sharing more personal information than what we do. It would be something similar to living in a small village (Armstrong 2013b). One might think that when everyone knows each other's secrets, no one feels particularly vulnerable. Privacy as we think of it, with its secluded refuges and moments of voluntary anonymity, is quite a recent innovation, says Brin (1998, 76). Maybe we do not need as much privacy as we have grown accustomed to.

Some of the ways we enjoy privacy may be relatively new, but our desire for privacy is as old as we are, as we saw in Chapter One. We crave privacy and seize opportunities for intimacy and solitude when they are available. That we had less opportunities for privacy in the past does not imply that our lives would be just as good if we went back to how things used to be. The lack of privacy in small villages can be suffocating, and introverts, eccentrics, and outliers suffer from it more than others. Mutual gossip does not seem to eliminate embarrassment, humiliation, shame, and stigma. The scrutiny one must endure in small villages can breed conformity. A lack of privacy also ties people to their past: villagers will never forget one's mistakes. In contexts where everyone knows too much about everyone else, reinventing oneself is hard.

Transparency advocates seem to be under the impression that, if everyone's personal information can be accessed, people will not be judged harshly anymore. After all, if you make me feel bad about my mistakes, I will make you feel bad about yours; it is better that we leave each other alone. Mutual transparency, however, is no guarantee against discrimination and unfair judgments. In most

cases of face-to-face interactions, race is something apparent. But transparency with regards to race has not made racial discrimination disappear, and I do not see how other kinds of discrimination and harsh judgments would evaporate with radical transparency.

So far I have talked about individuals' privacy. In an ideal society, there might also be a place for organisational privacy, however. According to Alan Westin,

> Just as individuals need privacy to obtain release from playing social roles and to engage in permissible deviations from social norms, so organizations need internal privacy to conduct their affairs without having to keep up a "public face." (Westin 1970, 44)

Without privacy, says Westin, the functioning of organisations would be impaired. Westin's assertion is too broad: we do not want to give organisations so much privacy that they have the freedom to engage in corruption. At the same time, it may well be beneficial for everyone if organisations can make *some* decisions in private.

Take the European Council of Ministers as an example (Stasagage 2006). Secret meetings and decisions at the Council carry important risks—that representatives might pursue private interests rather than public ones without the public being able to observe them, and that they might express one view to their constituents and the contrary view at the Council. If representatives tell one thing to their constituents but act differently at the Council, the public might vote for them only because they are ignorant of their actions and deceived by their words. The risks

of meetings and decisions being public, however, may be just as weighty, if not more so. Public meetings and votes can become displays, rather than instances of genuine deliberation and political negotiation. When representatives know their voters are watching, they might take excessively tough bargaining positions that demonstrate loyalty to their constituencies but are unsurpassable obstacles to deliberation and negotiation.

If we take the European Council to be a supranational institution that has as its main objective to reach compromises and agreements between European nations that will benefit European people, the functioning of the organisation may be more impaired by the inability to negotiate and deliberate (the risk of transparency) than by the risks of secrecy.[6] The Council itself admitted the following:

> The council normally works through a process of negotiation and compromise, in the course of which its members freely express their national preoccupations and positions. If agreement is to be reached, they will frequently be called upon to move from those positions, perhaps to the extent of abandoning their national instructions on a particular point or points. This process, *vital to the adoption of Community legislation*, would be compromised if delegations were constantly mindful of the fact that the positions they were taking, as recorded in Council minutes, could at any time be made public through the granting of access to these documents, independently of a positive council decision (Statement of Defence of the Council of the European Union in Case T-194/94, Brussels, 13 July 1994, 23-24).

[6] As a recent example, in October 2015, an article from *The Guardian* claimed that what was stalling negotiations between the European Union and the United Kingdom were fears of leaks: 'Downing Street has been determined to avoid putting its specific demands on paper for fear of leaks that' would leave 'Cameron vulnerable to becoming a hostage to those in the Conservative anti-EU ranks who will mock his shopping list as deficient' (Traynor 2015).

If public meetings become a show, there is a risk that the 'real' decisions will be made in backroom discussions, perhaps over lunch (Stasagage 2006, 167). If lunch is also public, then perhaps the most important bargaining moments will happen in bathroom stalls.

Another concern with meetings and votes not being public is that improper pressure may be put on public institutions by persons or lobbies learning prematurely that certain actions are proposed (Westin 1967, 48). Even with radical transparency in place, improper pressures would still be a worry. Lobbies and governments are not always scrupulous enough to be deterred by shame. It is well-known how big corporations such as Amazon and Google lobby their way out of paying taxes (Wheelwright 2016), and yet this public knowledge does not seem to be enough to deter such companies.

Finally, organisational privacy also protects individual privacy. In a radically transparent world, hospitals would disclose their patients' conditions, schools would unveil students' records, etc.

In conclusion, a radically transparent society does not seem to be better than a less transparent one for protecting privacy. A transparent society might still be desirable, however, if its benefits are so many that they outweigh the sacrifice in privacy.

Other problems and challenges of a radically transparent society: conformity, power imbalances, and dark spots

The biggest problem in a radically transparent society has already been mentioned à propos of privacy: conformity. With radical transparency we can expect a decrease in crime, but also a decrease in human variation and creativity, in bold ideas, in dissenters and activists. Humans tend to conform more when they are under the constant gaze of others.[7]

Many psychology experiments show how the presence of others can lead to conformity, but the Asch study is a classic one. Solomon Asch (1951) conducted his conformity experiment with groups of male college students who were asked to participate in a perceptual task. All but one of the participants was in fact a subject—unbeknownst to him, the rest were confederates. The actual objective of the study was to observe how the subject would react to the confederates' behaviour. Students got shown two cards: one with a line on it and the other with three lines labelled 'A,' 'B,' and 'C.' One of those three lines was the same length as the one on the first card, and the other two lines were very clearly longer or shorter. Participants had to identify the line that matched the one on the first card. In a control group where confederates did not pressure the subject, participants had an error rate of less than 1%. However, when participants were asked to say their answer aloud, and all confederates chose an incorrect line, the error rate

[7] This point can be seen as a modern version of John Stuart Mill's argument for a protected private sphere, partly to encourage experiments in living, individuality, and the freedom to form one's own opinion (Mill 1978).

increased to 36.8%. What this and other follow-up experiments show is that others can exert an enormous pressure to conform to popular views.

Another challenge for transparency is equality. Transparency advocates think that a greater equality will result from adding sousveillance to surveillance and coveillance. Equality, however, will not come for free. We already know (roughly) how much the richest 1% of the population earns. Knowledge is not enough to bring about change. People will have to protest and demand change. In a society that suddenly goes transparent but still suffers from grave inequalities, it is not at all clear that activists will not be persecuted by the powerful. Yes, the powerful would also be surveilled, but in practice, that might not deter them from exercising their power to ensure their privileged place in society before equality is secured.

A transparent society might work better if we were starting out from scratch. Given the highly unequal world we are stuck with as a starting point, I find it hard to believe that the mighty will, first, accept transparency for themselves. And, second, if radical transparency were somehow instantiated, that they would not take advantage of it to increase their power and squelch demands for more equality. Even Mann, Nolan, and Wellman admit the risk of exacerbating power imbalances:

> [T]he ubiquitous total surveillance that sousveillance (…) affords is an ultimate act of acquiescence on the part of the individual. Universal surveillance/sousveillance may, in the end, only serve the ends of the existing dominant power structure. Universal sur/sousveillance may support the power structures by fostering

> broad accessibility of monitoring and ubiquitous data collection. (Mann, Nolan, and Wellman 2003, 347)

If that risk were to materialise, there would be no decrease in police force and powers. Nor would there be fewer laws. On the contrary: the mighty could take advantage of the excess of unenforceable laws to make the legal system arbitrary. John Baker, a retired Louisiana State University law professor believes it is not an exaggeration to say that '[t]here is no one in the United States over the age of 18 who cannot be indicted for some federal crime' (cited by Fields and Emshwiller 2011). If the mighty were to enjoy total surveillance, anyone deemed uncomfortable to them could be indicted. Sousveillance might help expose injustice, but exposing something wrong is a long way from making it right.

In 2011, for example, the New York Police Department arrested Occupy Wall Street protestors for wearing masks in violation of a law from 1845 that bans two or more participants in a gathering from doing so (Robbins 2011). That law had probably not been enforced in decades. Even today, if you and a couple of friends stroll through Fifth Avenue wearing masks, the chances of getting arrested are next to none. There were no transparency issues in the Occupy case. Everyone understood that the law was being used unfairly as an excuse to arrest people who had become uncomfortable to the government. The protesters knew it and newspapers made it known to the wider public. It is a case that shows how transparency and injustice can coexist.

For starters, taking for granted that complete transparency will be achieved seems

like a big leap of faith. Brin argues that we cannot trust privacy laws because governments and businesses can simply spy on us behind our backs. He does not explain, however, how we can trust governments and businesses to carry all their operations out in the open. It is unclear whether transparency can be guaranteed through technical means. It might be as hard to police transparency laws as it is hard to police privacy laws. If we allow a sliver of bedroom privacy (as Brin wants us to), we cannot be sure government officials will not use their bunker bedrooms to keep a second set of books, or that tunnels will not be built from one bunker bedroom to another in order to carry out illegal meetings. Parallel underground systems of illegality could be built in office buildings and disguised as bathrooms (or similar excuses can be used to keep dark places out of the transparent society). A transparent society would make corruption a harder business to engage in, but people with financial resources would surely find ways to step out of the spotlight.

In what follows I conclude the chapter by balancing the benefits of a transparent society against the risks and downsides.

IV. Should we strive towards a radically transparent society?

There is no doubt that a radically transparent society would have many advantages. There would be research advancements in medicine, sociology, economics, and more. At least some global catastrophic risks (i.e., pandemics, weapons of mass destruction) would be prevented or mitigated. A transparent society would also be a convenient one, where walls, passwords and wallets

become unnecessary. Perhaps more importantly, a transparent society would be a safe one, as crime would be greatly reduced.

It is possible, however, that we may gain most of those advantages without having to give in to a radically transparent society. Many societies that are not radically transparent have nonetheless achieved a level of safety that is probably considered good enough by most people. Japan, Iceland, and Spain are some examples where less than one person dies a violent death per 100,000 citizens per year. Even radical transparency would not be able to stop some unpremeditated crimes. Similarly, we could get the research benefits of big data by anonymising personal information and controlling its access through, for example, differential privacy (a black box system that does not give up all its data to scientists, yet it allows researchers to ask specific queries and returns answers with mathematical noise to avoid the possibility of identifying individual subjects). Targeted surveillance when there is reason for suspicion is likely to be enough to minimise global catastrophic risks.

Radical transparency carries with it significant risks and disadvantages. It is unclear that it could be achieved; we may not be able to be certain that the powerful will not build parallel underground systems to escape sunlight. If we fail to achieve complete transparency, we risk worsening power imbalances and abuses through the expansion of mass surveillance on ordinary people. If we do succeed in achieving full transparency, most privacy will be lost, with all the psychological drawbacks that would entail, and the disadvantages for

organisations. We may be able to maintain secure bunker bedrooms, but even this is uncertain and will depend on technology to detect intrusions. With no privacy, a conformist society is welcomed where the risk of supporting unjust power structures might be high.

When advocating for more transparency, Justice Louis Brandeis is often quoted noting that '[s]unlight is said to be the best of disinfectants' (1913, 10). Most people forget that Brandeis was also one of the first two authors to advocate for a right to privacy, as we saw in Chapter Three. Sunlight can indeed disinfect, but in excess it can also kill life forms that we wish to preserve.

The risks and disadvantages of a transparent society come because of an excess in transparency—the world is not a better place if sunlight shines *everywhere* and *at all times*—and because transparency *by itself* cannot bring justice, equality, accountability, or safety. As O'Neill points out, transparency offers 'fewer and more limited benefits than is widely assumed' (2006, 89).

Transparency, then, needs a reliable rule of law, and robust systems of checks and balances, which include ways for people to challenge decisions and actions that have been made by others. The irony is that, once these practices are in place, we cease to need *radical* transparency to bring about desirable goals. Some degree of transparency is almost certainly needed to support accountability, but such transparency need not be radical—it does not have to be carried out in real time, it does not have to include *all* actions by individuals and organisations, and

whatever information is made accessible does not need to be made accessible to *everyone*.

Let us go back to Brin's two cities: the one where the cameras report to the police, and the one where the cameras report to anyone wanting to watch. We do not need to choose between them. We can construct cities without cameras; or where footage from cameras gets erased often and can only be accessed when it is necessary to solve a crime that has been committed or with an individual warrant when there is justified suspicion that a crime will be committed. However, if we had to choose between the two cities, given all the considerations we have gone through, the first city seems to be superior to the second—as long as it is set in a democratic country where police are held accountable. If there is no choice but to be watched, it is preferable to be watched by a limited group of people who are supervised and have to report to people who look out for citizens' interests (e.g., elected officials, a committee of representative citizens, and an ombudsman) than be watched by just anyone and potentially everyone. We should probably strive for a more transparent society, but not for a radically transparent one.

The crucial question, then, is when is transparency appropriate. Transparency is mostly inappropriate when it comes to individuals. As private citizens, we do not owe each other transparency. The idea that the more transparent a person is, the more virtuous she must be, is one that can only encourage superficiality and compliance, and that is bound to favour government power and corporations that profit from personal data. As individuals, it is enough that we can show that we

pay our taxes and are not free riders to secure a fellow feeling of trust and justice among citizens. Individuals who hold public offices may be amenable to more transparency requirements in virtue both of their salaries being paid by all citizens, and the power they hold to change the rules of society. However, they should only be made to disclose facts that are *directly* relevant to their public roles. For example, unless there is criminal activity involved, conflict of interest, or a direct relation to their work as public officials (e.g., if the politician involved is trying to regulate sexual lives in a way that contradicts his actions), the sex lives of politicians should remain a private matter.

Transparency, then, is mostly appropriate in institutional settings. Even then, limits are in order. First, given the costs of transparency, institutions should be asked to comply only with the minimum amount of transparency necessary to ensure good practices. Transparency can be limited by restricting both the information revealed and the audience that can have access to that information, as well as by delaying revelations in time. Anything above the minimum necessary should not be mandatory. Second, as far as possible, transparency should be about processes and practices, rather than about the individuals participating in institutions. Third, when the costs of transparency outweigh the risks of secrecy (e.g., by severely impairing the functioning of an organisation, by creating more opacity, or by having a severe negative effect on people's wellbeing), transparency should be further limited.

CHAPTER SEVEN

Privacy and the Moral Dangers of Decision-making Algorithms

More and more of our personal information is being collected and analysed. This fact may seem like a significant threat to privacy. It is increasingly the case, however, that humans may never look at that information. Instead, algorithms are performing privacy-invasive analyses and making decisions on the basis of them. An algorithm is a series of instructions that tell a computer what to do (Domingos 2015, 1). Whether people are aware of it or not, algorithms are increasingly ruling our lives. They filter candidates for employers, assess loan applications for banks; match people on dating sites, recommend products on behalf of online businesses, assess both teachers and students in educational institutions, buy and sell stocks (Pasquale 2015, 129), analyse citizens and win over their votes for political campaigns (Grassegger and Krogerus 2017), look for terrorists, decide what to show you on social media and search engines, and much more. There are fewer and fewer areas of life that are not heavily influenced by algorithms. According to data scientist Pedro Domingos, 'If every algorithm suddenly stopped working, it would be the end of the world as we know it' (2015, 1).

It might be comforting to think that, given that we are living in a world in which much of our personal data is being collected and analysed, it is better for privacy to have algorithms sifting through our data than it is to have humans do it. After all, nobody would blush at the thought of an algorithm 'knowing' facts about our private life. Furthermore, algorithms can process data much faster than humans, promising time saving and effectiveness, and they can, in principle, avoid human errors such as implicit bias. Algorithms, however, also have a dark side. They too can be prone to biases and can lead to serious injustices. And they are powerful tools that enable privacy losses and violations of the right to privacy.

Section I discusses why we cannot lose privacy to algorithms and why they cannot violate our right to privacy. This does not mean, however, that algorithms are free from moral problems. In Section II, I go through some problems that algorithms can fall prey to, such as bias and discrimination. In Section III, I describe the current opacity of decision-making algorithms and briefly explore some possible solutions. Section IV concludes by making some final remarks about privacy and algorithms.

I. We cannot lose privacy to algorithms, and they cannot violate our right to privacy

If algorithms were instantiated by people and not computers—if people were the ones sifting through our most intimate data—our right to privacy would, without a doubt, be violated, and we would be losing much privacy. As things stand,

however, computers are often in charge of going through our personal information. In Chapter Four we saw that privacy losses usually lead to feelings of embarrassment and self-consciousness. Other things being equal (supposing there are no further consequences), do you care about an algorithm accessing information about your sex life, or your medical history? Probably not.

The feeling that, from the point of view of privacy, it is more worrisome to have a human look at our private information than to have an algorithm do it comes, I suspect, from the thought that algorithms are not our peers. Computer algorithms do not care about intimate details about your life the way a human might—they are not sentient beings and they do not have an understanding of social life like we do. Algorithms do not judge; they do not get excited at the sight of a naked body; they do not feel disgust or outrage; they do not stare; they do not gossip; they do not laugh at our faults.

Similarly, algorithms cannot violate our right to privacy because they are not moral agents. An agent is, roughly, an entity that can be the source of action. The requirements for being a *moral* agent are not uncontroversial, but algorithms do not make the grade under most plausible criteria. Gary Watson (2013) convincingly argues that moral agents are beings that are autonomous (self-governing), and accountable (answerable to others). In what follows I will argue that algorithms are neither.

Luciano Floridi and J.W. Sanders contend that algorithms are autonomous because they are 'able to change state without direct response to interaction'—that is, they can act independently of the humans who created them (2004, 357). Autonomy in moral and political philosophy, however, is a much richer concept.

Crucial to autonomy is the capacity to act in accordance with reason in a way that responds to one's own motives (Christman 2015). To be autonomous, one must be able to reflect on (Watson 2013, 4-5), endorse, and act on one's values (Christman 2015). It is because a person is able to choose her values for herself and live accordingly that we must ask for her consent to interact with her in invasive ways (for example, in the case of a medical procedure). We do not need to ask an algorithm its permission to modify it or even terminate it because algorithms do not have values of their own. Nor do they seem to respond to reasons qua reasons.

Algorithms are programmed to do something: win a game of chess, distinguish spam from non-spam, find people who might be interested in buying a product, assess whether a candidate will be appropriate for a job description, etc. At the moment, however, algorithms are incapable of normatively assessing the objective for which they have been created and modifying their behaviour accordingly. Consider the role algorithms play in advancing for-profit colleges in the United States. These are expensive, low-quality colleges that advertise themselves to vulnerable populations as a way out of their underprivileged status. In fact, in the work market, a person is no better off having a diploma from a for-profit college than having not attended college at all (Darolia et al. 2014). When an algorithm is

looking for a possible client for a for-profit college, it looks for people in the poorest postal codes who have clicked on ads for payday loans or whose search histories show a concern with post-traumatic stress (O'Neil 2016, Loc 1052). When algorithms do their tasks, they are not wondering whether it is morally correct to prey on such people, and they are incapable of deciding to quit their jobs and go for a more ethical line of work.

As Domingos puts it, 'computers don't have a will of their own' (2015, 283). They do not reflect on what they want, on what is worth pursuing, or on how they should live their lives. Admittedly, there might come a time when artificial intelligence becomes so sophisticated that computers will appear to exhibit and perhaps will even possess desires and values of their own. At the moment, however, I see no evidence to think this is the case.

Just as they are not autonomous, algorithms are not accountable. As accountable beings, 'we are answerable to others for how we lead our lives' (Watson 2013, 1). That is, we can recognise others' interests and moral claims, and when we do not respect them, we are liable to be the subjects of complaints or even punishment. An algorithm, in contrast, does not think about the suffering it might be causing by encouraging vulnerable people to take out heavy loans in order to pay for a degree at a for-profit college that is worth little or nothing. Similarly, when wronged by an algorithm, it would not occur to us to punish it or ask it for compensation. Rather, we would seek redress from the people who designed and implemented the algorithm.

Floridi and Sanders (2004) argue that we should not confuse accountability with responsibility. According to them, '[a]n agent is morally accountable for *x* if the agent is the source of *x*,' where *x* is an action causing moral good or evil. To also be morally responsible, they argue, 'the agent needs to show the right intentional states' (Floridi and Sanders 2004, 371). They believe that entwining the concepts of accountability and responsibility amounts to 'confusing the *identification* of *x* as a moral agent with the *evaluation* of *x* as a morally responsible agent' (367). Morality is intrinsically about normative evaluation, however. If a moral agent can be identified as such, then it must also be the case that we can evaluate her as responsible for her actions.

Floridi and Sanders believe that there is such a thing as 'mind-less morality' (2004, 351). I disagree. I find it quite likely that mind and sentience are necessary for moral agency. In order to have a conception of the good that we want to pursue (autonomy), we need to have a feel for what leads to pleasure, meaningfulness, and satisfaction; and in order to guide one's actions by the recognition of others' moral claims in a way that can count as a moral action (accountability), one must have a sense of others' capacity to suffer, of what it feels like to be harmed. Perhaps a computer scientist could programme an algorithm to behave in such a way as to not make people frown (a proxy for not making people suffer), thereby roughly behaving in a moral way. Such an algorithm may have moral consequences, but it would not be a moral agent because it would not be acting from moral reasons, but from a set of instructions. This is the moral equivalent of the hard problem of

consciousness. We could thus ask whether moral zombies—beings who act indistinguishably from normal moral agents but for whom there is nothing it is like to be them; beings who do not feel suffering, empathy, regret, or anything else—could count as moral agents. I think not.

That we do not lose privacy to algorithms and that they do not violate our right to privacy does not mean that they cannot cause moral havoc. In the next section I go through some cases that exemplify the dangers of decision-making algorithms.

II. Algorithmic disasters

Everyone knows that humans can be biased. Most, if not all of us, have been both victims and agents of biases. In contrast, there seems to be a blind faith in the objectivity of computers and what they can do (Danielle Citron, cited by Dormehl 2014). We forget, however, that algorithms learn from data that is, directly or indirectly, collected by humans, and they process that data with instructions programmed by humans. Far from being impartial, models and algorithms reflect the judgments, priorities, and biases of their designers (O'Neil 2016, Loc 297). What follows is a brief exploration of cases of algorithmic biases and moral failings.

Researchers from the University of Washington found that if you search in Google Images for pictures of CEOs, only 11% of those will be women (cited by Miller 2015). This does not even reflect the already worrisome proportion of women who

are chief executives in the United States (27%). Similarly, in a recent study, researchers disclosed how Google showed ads for high-paying jobs more often to men than it did to women (Datta, Tschantz, and Datta 2015). In both cases, algorithms are putting women at a disadvantage: first, by reinforcing a sexist paradigm in which positions of power are held mostly by men; second, by acting like a self-fulfilling prophecy through limiting women's possibilities of attaining a high-income job by not showing them the corresponding ads.

Worse still, Google's autocomplete feature has been known to reinforce even more alarming stereotypes regarding women. In 2013, the UN developed an ad campaign revealing autocomplete results for searches such as 'Women should…' (stay at home, be slaves, be in the kitchen, not speak in church), 'Women shouldn't…' (have rights), and 'Women need to…' (be disciplined) (Mahdawi 2013). In this case, the algorithm may simply be reflecting what people actually search for, what people believe, but it is still playing a part in reinforcing these beliefs, as people who would have perhaps searched for something else can suddenly find themselves being influenced by the first options presented. (Since 2013, Google has modified its algorithm to get rid of such dreadful autocomplete suggestions about women.)

Similar prejudiced behaviours on the part of algorithms have been observed towards blacks. Latanya Sweeney (2013) discovered that Google ads were significantly more likely to show ads for arrest records on searches for names that are typically assigned to black people.

The algorithms mentioned so far are not innocuous: they are having a harmful influence on our culture, on preconceived ideas about women and blacks, and in some cases they are unfairly limiting people's access to opportunities. But harm created by algorithms can be even more tangible and egregious. In some cases, algorithms are responsible for people being denied jobs unfairly, getting fired, not getting loans (or getting very expensive loans)—any of which can throw a person into a vicious cycle of poverty and disadvantage.

In some cases algorithmic unfairness is the result of having questionable or false models or proxies. Consider algorithms that calculate the risk of recidivism in criminals in order to guide sentences. One of the most used models, the LSI-R, uses a questionnaire that includes questions about the first time the prisoner was 'involved with the police,' and whether the prisoner has friends and relatives with criminal records (O'Neil 2016, Loc 357). The answers to these questions will further disadvantage people who already come from underprivileged neighbourhoods. Early involvement with the police, for example, could be related to a racist stop and frisk for which the prisoner is hardly at fault. The justice system is supposed to judge people on their actions, but algorithms that take into account circumstances that the criminal is not responsible for, judge people for who they are (rather than what they have done) (O'Neil 2016, Loc 383), and on the basis of what people similar to them have done in the past.

Other examples of algorithms that use inappropriate models can be found in the field of education. Teachers have been wrongfully rated—and hired, sacked, or warned accordingly—by algorithms based on models that were later shown to be faulty (O'Neil 2016, Ch 1, Ch 7). In one case, blogger Gary Rubinstein found that about a fourth of teachers had a 40-point difference in scores in consecutive years and teaching the same subject, which made him suspect the scoring algorithm was not tracking teachers' abilities, which would likely have been more stable (O'Neil 2016, Loc 1946). In fact, the algorithm was rating teachers on the basis of how well students fared compared with forecast models. The algorithm measured the gap between how well it expected students to do and how well students actually did. But the forecast models were highly speculative, derived themselves from statistics, with so much mathematical noise as a result, that in the end it turned out the algorithm was measuring nothing at all (O'Neil 2016, Loc 1927). The scoring system led to utter arbitrariness.

Yet another sphere of life where algorithms use questionable proxies is personal finance. Banks calculate how likely people are to pay back credit by using proxies that are far from ideal. In the United States, Facebook has patented a new type of credit rating based on the credit ratings of one's friends on social media (Meyer 2015). People who are Facebook friends with individuals with bad credit histories may be denied a loan. There is no evidence Facebook is using this system (yet), although it is not clear to me that people would be informed if it was being used. In any case, social media relations are already being used for credit purposes by other businesses. An example is the German company Kreditech, which asks loan

applicants to share information about their social media networks (Waddell 2016). As most people tend to associate with people of their own race, this kind of proxy can closely resemble racial discrimination, and in any case will surely disadvantage people belonging to poor communities. People should be judged on whether they will be able to pay back a loan on their own merits, and not on the relations they keep.

In some other kinds of cases, algorithmic moral failures come about as a result of misidentifying people. In one case, a man got his driver's licence revoked because a facial-recognition algorithm had confused him with another driver. When the man complained, he was told that, because protecting the public is weightier than the inconvenience of wrongly targeting a few people, the burden was on him to clear his name (Dormehl 2014).

In a similar vein but with graver consequences, a woman called Catherine Taylor was denied a job because she was mistaken for another woman with the same age and date of birth who had been charged for making and selling methamphetamines (Mui 2011). Correcting this mistake has proved impossible. The experience has destroyed her credit record, and she has been turned down for an apartment, as well as rejected as a volunteer for Girl Scouts. Not surprisingly, she believes the stress of her situation has had a negative impact on her health (Mui 2011). Correcting such mistakes is hard because there might be dozens of data brokers who are selling, or have already sold, the incorrect information to other interested parties. More importantly, in such cases the

victims are not clients but products being sold, so data brokers have no incentive to help them (beyond the satisfaction of doing the right thing) (O'Neil 2016, Loc 2128). Data brokers stand to earn nothing by taking the trouble of correcting a mistake that leaves someone out of the game of life. People like Taylor end up being collateral damage.

In yet other cases, algorithms damage lives on the basis of accurate information that nevertheless should not be taken into account. Helen Stokes had been arrested twice during fights with her husband, but those arrests did not legally exist—she had never been convicted and had successfully had the records legally expunged from government databases. Nevertheless, her records showed up for a corporation providing background checks on tenants, and as a result she was prevented from moving into a more affordable apartment (Palazzolo and Fields 2015). Arrests that do not result in conviction (particularly those that have been expunged by the appropriate authorities) should not be taken into account in a person's record because they are not indicative of any moral or legal fault—anyone can get wrongly arrested.

In a nutshell, algorithms can cause injustice by discriminating against people on the basis of criteria such as their postal code, the kinds of stores they buy in, and their interests—which are proxies for information such as their race, gender, and purchasing power. If treating people differently on the basis of these latter features is morally unacceptable because it is discriminatory, then so is relying on their proxies. Algorithms can also cause unfairness by being based on false models that

lead to arbitrariness, through mistakes such as the misidentification of people, and by taking into account information that should not inform their decisions.

In some cases, the fault may lie in the data being fed to the algorithm. In other cases, the fault lies with the proxies the algorithm is using. Proxies are often not programmed into algorithms. Rather, it is the algorithm itself that selects the proxy through a process of machine learning, such that proxies may be hard to predict or even detect. In yet other cases, we may never understand the cause of the unfair outcome. Algorithms identify correlations without establishing causation (Mittelstadt et al. 2016). Sometimes, the correlations that will then form the basis of proxies will be discriminative (as in the cases of postal codes, or Facebook friends). Other times, the correlation found may be spurious, tracking nothing.

In our current state of affairs, most cases of unfairness caused by algorithms will never be known—let alone corrected. Most of us will never question what we see on our Facebook feeds, or our Google searches, since we usually do not have the chance to compare what we see to what other, very different people, see. Most people who are denied a loan or a job do not get an explanation as to why and will not inquire further. And moral failures will never be flagged as failures because the benchmark used to measure the success of algorithms is usually profit, or efficiency. As long as the money keeps coming in, no one will notice that people are being treated unfairly. Even victims may not realise they are victims.

Algorithms are not—and, some argue, cannot be (O'Neil 2016, Loc 2182)—programmed to judge what is fair. Only humans can do that.

Algorithms, then, not only run into problems of fairness. Because more often than not we do not realise the moral havoc they are causing, they also suffer from a problem of opacity.

III. Opacity

If algorithms were faultless—if they measured reality and not proxies, if they gave each person what she deserves, if they did not work against the benefit of people—their opacity would not be as problematic. But algorithms can and do seriously harm people. Consequently, it is important that we be able to peer into their workings to determine what went wrong and how it can be fixed. For that, we have to understand algorithms better. It is easier said than done.

Algorithms are sets of instructions instantiated by computers. Computers are composed of billions of transistors—tiny switches that can be turned on and off. Bits of information are nothing but states of transistors: ones if the transistors are on, and zeros if they are off. The simplest algorithm instructs a computer to flip a transistor. The second simplest algorithm tells the computer to combine two bits (Domingos 2015, 1). In this chapter, I am only concerned with decision-making algorithms—that is, algorithms that do not need any human intervention to go

through data and make decisions that impact the lives of people in significant ways.

Throughout the history of computer science, algorithms have become increasingly complex: first, because new algorithms are built on top of old ones, creating a sort of algorithmic tower in which generations of algorithms interact in unexpected and complicated ways (Domingos 2015, 5); second, because programmers can now design learning algorithms—algorithms that create their own algorithms (Domingos 2015, 5). Machine learning enables computers to programme themselves in ways that computer scientists may no longer understand. In other words, computers are writing programs that people cannot write, because we do not know how to tell computers how to do certain things, such as how to decipher handwriting (Domingos 2015, 6). However, if we give a learning algorithm enough examples, it can train itself to read handwriting. Learning algorithms are thus creating programs that are millions of lines long—too big and complex for human understanding (Domingos 2015, 7). As a result, machine learning algorithms are hard to predict (i.e., know how inputs will be managed) and explain (i.e., understand how a decision has been made) (Mittelstadt et al. 2016, 3-4).

Complexity, however, is not the only cause for decision-making algorithms being black boxes. Companies wish to keep their algorithms secret to avoid competition. They argue that algorithms are intellectual property, and constitute a 'secret sauce' to their business, which in the case of Google, Amazon, and Facebook, makes them worth hundreds of billions of dollars (O'Neil 2016, Loc 413).

A further justification for keeping algorithms secret is to stop people from 'gaming the system' (Domingos 2015, 35). Algorithms work on models of the world. Models, by definition, are simplifications of complex realities (O'Neil 2016, Loc 288). If a model includes all possible variables, it is no longer useful, as it is as difficult to analyse as the world itself. To create a model, then, choices are made as to what is to be included (O'Neil 2016, 291). In ranking universities, for example, criteria such as SAT scores, student-teacher ratios, acceptance rates, graduation rates, and job prospects are typically taken into consideration as proxies for a good education (O'Neil 2016, Loc 726). Proxies, however, can only approximate what they purport to measure, and can often be manipulated. Instead of improving the education which they deliver, universities can focus on bettering the relevant proxies through questionable methods. By paying the fee for admitted students to retake the SAT, they can boost their score. To increase graduation rates, they can lower the standards for difficult subjects such as math (O'Neil 2016, Loc 931). Universities have also been known to hire their own graduates for hourly temporary jobs in order to boost their job placement numbers (Segal 2011). If the proxies used by algorithms are kept secret, people will have a harder time gaming the system.

Finally, institutions may have an interest in keeping their algorithms secret in order to shield themselves from possible criticism. If algorithms are out of view, people may not even know they exist, let alone question or challenge their workings.

The standard response to the problem of opacity is to push for more transparency. Angela Merkel, for example, has called for more algorithmic transparency (Connolly 2016). Having more transparency is likely to be a necessary (but not sufficient) condition to deal with algorithmic unfairness. It is particularly important to demand transparency with respect to the goals of decision-making algorithms. Algorithms are programmed to do *something*, to achieve a goal, and authorities should have knowledge of that goal in order to regulate it. In the case of predatory loans, for example, if algorithms are looking for desperate and uneducated people willing to pay absurd interest rates, then the goal of the algorithm must be modified into something morally acceptable.

Transparency, however, has its limits. In the case of machine learning, algorithms are likely going to be beyond the reach of human comprehension—even if we have access to every line of code. One possibility is to have governments ban any algorithm that cannot be fully explained to an expert auditor. On this view, people who commission and programme algorithms are responsible for their creation, and it would be illegal to create something that they will no longer be able to understand or fully control. The downside of this option, however, is that the capabilities of algorithms will be limited severely. Computer scientists and corporations would surely fight back, although it is perhaps the case that programmers will soon manage to design machine-learning algorithms that are able to explain themselves when asked about a particular decision.

A measure that is likely to be more successful in the short-run is the European regulation that decisions that impact people negatively should not be carried out solely by algorithms, and that people have a right to an explanation (Directive (EU) 2016/680 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). In other words, if someone is going to be denied a loan, or a job, for example, on the basis of an algorithm, a human should revise that decision, make sure it is a fair one, and explain the reasons for rejection to the person being denied a good. It is one thing to let algorithms process information and do work for us. But decision-makers who are moral agents must assume responsibility for harming people. If people have a right to an intelligible explanation for being denied an opportunity, then incomprehensible algorithms should not be used for making decisions that can place people at a significant disadvantage.

I will end with some final comments about privacy and algorithms.

IV. Conclusion: privacy and algorithms

In this chapter I argued that algorithms are not a problem for privacy, and then went on to explore other moral issues related to decision-making algorithms. Now that these have been reviewed, some final remarks about privacy are in order.

I stand by my contention that algorithms having access to sensitive data is not a privacy problem because algorithms are not our peers and are not moral agents.

Someone could argue, however, that I am wrong because the negative effects of algorithmic decision-making are the same as the negative effects of privacy losses. Consider a woman wanting to keep private her homosexuality, among other reasons, to avoid being discriminated against in a job application. Suppose she does get discriminated against on the basis of that information. If an algorithm rejected her, then it is only an issue of fairness, according to my view. If a human rejected her, then it is both a privacy and a fairness issue. If the cause (a piece of information) and the effects are the same, does it make sense to maintain the distinction?

I think it does. It is a distinction similar to that of privacy perceptions and privacy losses. Just as a reminder from Chapter Four: if someone puts a sham camera in your bedroom without your knowing it is a fake camera, it will have some of the same negative effects as if it were a functioning camera (e.g., it might make you feel uncomfortable, it might alter your behaviour, etc.). Nevertheless, it does not constitute a privacy loss because there is no one who knows more about you as a result of the sham camera, and there is no one actually accessing your autotopos. The *effects* of privacy losses (or other wrongs like the manipulation of privacy perceptions) should not be confused with the privacy loss itself. If the camera were real, you would also feel uncomfortable, and it might also alter your behaviour, but on top of it there would be the effects of the privacy loss: at the very least, the person having access to the footage would know more about you, and she would have access to your autotopos (with all the implications this carries about how she could use that information and access).

Admittedly, some of the effects of algorithms having access to sensitive data are the same as the effects of privacy losses: discrimination, financial losses, etc. However, these are only effects—they are not part of the privacy loss itself. Moreover, privacy losses carry extra effects that are not present in the case of algorithms having access to sensitive data. Let us return to the homosexuality case. In the case of the prospective employer who is a human, there is now a moral agent and a peer knowing something very sensitive about someone—he could make her feel embarrassed, he could have reason to share that information with other members of the community to harm that person, etc.

That we cannot lose privacy to algorithms is not to deny, however, that algorithmic decision-making is an important part of 'surveillance capitalism'—a term coined by Shoshana Zuboff (2015) to describe an economy fuelled by the monetisation of personal data acquired through surveillance. Algorithms are a crucial part of an array of tools that are putting our privacy at risk. They can impact privacy, first, by collecting data. Once the data is collected, it can be used against data subjects. Among other perils, it can be stolen and accessed by criminals, enemies, governments, etc.

Second, algorithms infer sensitive information from non-sensitive information. Think of all the information that is out there about you on the Internet, and try to imagine what a smart algorithm could infer from it. Some inferences are rather obvious. It is easy to understand that if you Facebook Like many businesses in

your local area, it will be straightforward to infer what city you live in. But because algorithms work with correlations, some associations can be utterly surprising. Cambridge researchers, for example, have discovered that Liking Curly Fries was one of the best predictors for high intelligence, even though liking fries does not seem to be causally related to being smart (Kosinski, Stillwell, and Graepel 2013, 5804).

Algorithms, then, do represent a threat to privacy, but only indirectly, by collecting data and inferring sensitive information. For this reason, it still makes sense to say that, for privacy reasons, we do well in shielding our information from prying algorithms. If our data is being collected, and if sensitive information is being inferred about us, our right to privacy is being violated, not by algorithms themselves, but by the people who access that data or put themselves in a position in which they can access that data. Algorithms enable privacy losses and violations of the right to privacy, but they are merely tools, not moral agents.

Conclusion

The aim of this dissertation has been to contribute to a better understanding of privacy. In Chapter One, the history of privacy was explored; it was suggested that privacy is a universal human need, likely grounded in our animal origins. The relationship between privacy and the private and public spheres was clarified in Chapter Two. There, I argued that privacy does not map onto the distinction between the private and the public, and that arguments need to be made independently of the appeal to that divide to defend what privacy is and what it ought to cover. Chapter Three reviewed the most common definitions of privacy in the academic literature, pointing out their shortcomings and acknowledging their respective strengths. Chapter Four offered a map of the moral territory of privacy. The map started with the definition of privacy as remaining personally unaccessed. I then explored the value of privacy, offered a defence of the right to privacy as a claim to a robustly demanding good, and clarified the role of social norms in the right to privacy. The chapter ended with an overview of the moral significance of perceptions of privacy, and a defence of the duty to protect one's own privacy.

The last part of the dissertation was dedicated to practical matters. Chapter Five dealt with the balance between privacy and security, and argued that mass surveillance is a disproportionate, unnecessary, and ineffective response to the

threat of terrorism. The chapter also defended the wide use of encryption, as it is a method that constrains the mass surveillance of data without seriously compromising authorities' abilities to fight crime. Chapter Six explored the balance between privacy and transparency, and argued against having a radically transparent society. Chapter Seven delved into the issue of privacy-invasive algorithms. I argued that algorithms risk our privacy by collecting data and inferring sensitive information from aggregated data. The moral responsibility for the collection and management of sensitive data, however, lies with moral agents; that is, with the people who programme and use privacy-invasive algorithms.

At least three broad lessons can be discerned from these explorations. The first is that privacy is more fundamental to society than it can appear: it is as old as humanity, universal, deeply entrenched in social norms, and vital to the psychological wellbeing of individuals and the political health of democracies.

The second lesson is that there is much at stake in losses of privacy. Looking at someone is a way of exercising power over him. One exercises power through looking itself, and not only through learning more about someone and potentially using that information against him. The act of looking is not neutral: it burdens people and pressures them into conformity. We may have come a long way from the savannahs we evolved in, but it is still the case that the stare of another turns us into potential prey.

Violations of the right to privacy leave us vulnerable in three ways. First, we become vulnerable to other individuals, from cybercriminals, stalkers, and trolls, to abusive ex-partners and personal or professional enemies. Potential harms include financial loss, identity theft, online harassment, physical insecurity, and public humiliation. Second, violations of the right to privacy make us vulnerable to private corporations that benefit from our personal data. The harms thus incurred can include unjustly being denied a job, an apartment, or a loan, being charged more than what is fair, having employers encroach on the personal lives of employees, and more. Third, violations of the right to privacy leave us vulnerable to governments and intelligence agencies. As opposed to justified infringements of the right to privacy, which subject criminals to a fair rule of law, violations of the right to privacy encourage authoritarian tendencies. They enable authorities to harass political opponents, and they risk the fabric of democracies. In a worst-case scenario, if a system of mass surveillance with detailed sensitive information collected from everyone for decades were to fall into the wrong hands, it could invest a tyrannical regime with power the likes of which have never been seen before. The collection of so much sensitive data is putting us all at risk, as the safety and good use of that data can never be completely guaranteed.

The third lesson is that, in the balance between privacy and other values such as security, transparency, convenience, or efficiency, privacy is not necessarily the weakling value, the dispensable luxury, or the necessary sacrifice. There are ways of achieving security that do not necessitate the surrendering of privacy. Transparency is not an absolute value that must trump all claims to privacy.

Convenience and efficiency may not be worth the losses of privacy we are suffering. It is not only that we ourselves might be better off safeguarding our privacy—we also have duties to protect privacy for the sake of others.

We are at a crossroads—a historical moment where important decisions are being made that will determine the role of personal data in our society for decades to come. It might be the case that we do not want to live in a world where every word we say and write, every interaction with others, every movement, every expression of love or friendship, is recorded. Perhaps personal information, like votes, is not the kind of thing that should be bought and sold. It may be a good time to explore alternative ways of funding the Internet. Maybe the police and intelligence agencies do not need to collect as much sensitive information from all citizens in order to keep us safe.

We can use encryption. We can demand privacy from corporations, and choose privacy-respectful services. We can refuse to give up our personal information when companies ask for it needlessly. We can put pressure on our politicians to make wiser choices. Perhaps we can safeguard bubbles of privacy where relaxation, creativity, intimacy, freedom of thought, speech, and association, can enable us to thrive. Perhaps we can build a better tomorrow. Let us not surrender the ability to write a diary, read an incendiary book, research a sensitive topic, talk to our doctors, lawyers, and friends, make fools of ourselves, give a kiss, or say 'I love you' with the peace of mind that comes from knowing that no one is watching.

Bibliography

2008. Protecting Individual Privacy in the Struggle Against Terrorism. Washington, D.C.: National Academy of Sciences.

2009. Report on the President's Surveillance Program.

Ackerman, Spencer. 2014. "FBI director attacks tech companies for embracing new modes of encryption." *The Guardian*, 16 October 2014. https://www.theguardian.com/us-news/2014/oct/16/fbi-director-attacks-tech-companies-encryption.

Allen, Anita. 1988. *Uneasy Access: Privacy for Women in a Free Society*. New Jersey: Rowman and Littlefield.

Allen, Anita. 2011. *Unpopular Privacy*. Oxford: Oxford University Press.

Allen, Anita. 2013. "An Ethical Duty to Protect One's Own Information Privacy?" *Alabama Law Review* 64 (4):845-866.

Andrews, Richard V. 1979. "The Physiology of Crowding." *Comparative Biochemistry and Physiology* 63A:1-6.

Angwin, Julia. 2014. *Dragnet Nation*. New York: Times Books.

Arendt, Hannah. 1998. *The Human Condition*. Chicago: University of Chicago Press.

Aristotle. 2013. *Aristotle's Politics*. Translated by Carnes Lord. Second Edition ed. Chicago: The University of Chicago Press.

Armstrong, Stuart. 2013a. "How to get positive surveillance—a few ideas." *University of Oxford Practical Ethics Blog. Ethics in the News.*

Armstrong, Stuart. 2013b. "Life in the Fishbowl." *Aeon*, 30 September 2013.

Asch, Solomon E. 1951. "Effects of group pressure on the modification and distortion of judgments." In *Groups, leadership and men*, edited by H. Guetzkow, 177-190. Pittsburgh, PA: Carnegie Press.

Ball, James. 2013. "NSA monitored calls of 35 world leaders after US official handed over contacts." *The Guardian*, 25 October 2013. https://www.theguardian.com/world/2013/oct/24/nsa-surveillance-world-leaders-calls.

Beardsley, Elizabeth. 1971. "Privacy: Autonomy and Self-Disclosure." In *Privacy: Nomos XIII*, edited by J. Roland Pennock and John W. Chapman, 56-70. New York: Atherton Press.

Benn, Stanley. 1971. "Privacy, freedom, and respect for persons." In *Nomos XIII: Privacy*, edited by J. Roland Pennock and John W. Chapman. New York: Atherton Press.

Bentham, Jeremy. 1995. "Panopticon or The Inspection-House." In *Jeremy Bentham: The Panopticon Writings*, edited by Miran Božovič. London: Verso.

Bentham, Jeremy. 2001. *Writings on the Poor Laws*. Vol. 1, *Quinn, Michael*. Oxford: Oxford University Press.

Bergen, Peter, David Sterman, Emily Schneider, and Bailey Cahall. 2014. Do NSA's Bulk Surveillance Programs Stop Terrorists?

Bergman, Lowell, Eric Lichtblau, Scott Shane, and Don van Natta Jr. 2006. "Domestic surveillance: the program; spy agency data after Sept 11 led F.B.I. to dead ends." *The New York Times*, 17 January 2006. http://www.nytimes.com/2006/01/17/us/front page/domestic-surveillance-the-program-spy-agency-data-after-sept.html.

Bertrand, Natasha. 2015. "The FBI staged a lovers' fight to catch the kingpin of the web's biggest illegal drug marketplace." *Business Insider UK*, 22 January 2015. http://uk.businessinsider.com/the-arrest-of-silk-road-mastermind-ross-ulbricht-2015-1.

Bezanson, Randall P. 1992. "The Right to Privacy Revisited: Privacy, News, and Social Change, 1890-1990." *California Law Review* 80 (5):1133-1175.

Birdwell, Jonathan, Charlie Cadywould, and Louis Reynolds. 2014. Tune in, turn out. Demos.

Bloustein, Edward J. 1978. *Individual and Group Privacy*. New Brunswick: Transaction Books.

Brandeis, Louis D. 1913. "What Publicity Can Do." *Harper's Weekly*, 20 December 1913, 10-13.

Brin, David. 1998. *The Transparent Society. Will Technology Force Us to Choose Between Privacy and Freedom?* Cambridge: Perseus Books.

Brown, Ian, and Douwe Korff. 2009. "Terrorism and proportionality of Internet surveillance." *European Journal of Criminology* 6 (2):119-134.

Brown, Jerram L., and Gordon H. Orians. 1970. "Spacing Patterns in Mobile Animals." *Annual Review of Ecology and Systematics* 1:239-262.

Brown, Peter. 1987. "Late Antiquity." In *A History of Private Life. From Pagan Rome to Byzantium*, edited by Paul Veyne, 235-312. Cambridge: Harvard University Press.

Bryant, Ben. 2016. "VICE News Investigation Finds Signs of Secret Phone Suveillance Across London." *VICE*, 14 January 2016. https://news.vice.com/article/vice-news-investigation-finds-signs-of-secret-phone-surveillance-across-london.

Bryson, Bill. 2010. *At Home. A Short History of Private Life*. London: Transworld Publishers.

Buijs, Stephanie, Linda J. Keeling, Carl Vangestel, Jeroen Baert, Jürgen Vangeyte, and Frank A. M. Tuyttens. 2010. "Resting or hiding? Why broiler chickens stay near walls and how density affects this." *Applied Animal Behaviour Science* 124 (3-4):97-103.

Chomsky, Noam. 2016. A Conversation on Privacy With Edward Snowden, Noam Chomsky, and Glenn Greenwald. In *The Intercept*.

Christian, John J., Vagn Flyger, and David E. Davis. 1960. "Factors in the Mass Mortality of a Herd of Sika Deer, *Cervus nippon*." *Chesapeake Science* 1 (2):79-95.

Christman, John. 2015. Autonomy in Moral and Political Philosophy. In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta.

Cicero, Marcus Tullius. 2009. *The Fourteen Orations (Philippics) of Cicero against Marcus Atonius*. Translated by C. D. Yonge: Digireads.com.

Clarke, Richard A., Michael J. Morell, Geoffrey R. Stone, Cass R. Sunstein, and Peter Swire. 2013. Liberty and Security in a Changing World. Report and

Recommendations of The President's Review Group on Intelligence and Communications Technologies.
Confucius. 1938. *Analects*. Translated by Arthur Waley. London: G. Allen and Unwin.
Connolly, Kate. 2016. "Angela Merkel: internet search engines are 'distorting perception'." *The Guardian*. https://www.theguardian.com/world/2016/oct/27/angela-merkel-internet-search-engines-are-distorting-our-perception.
Crisp, Roger. 2013. Well-Being. In *Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta.
Crovitz, L. Gordon. 2009. "Transparency Is More Powerful Than Regulation." *The Wall Street Journal*, 30 March 2009. http://www.wsj.com/articles/SB123837223623167841.
Darolia, Rajeev, Cory Koedel, Paco Martorell, Katie Wilson, and Francisco Perez-Arce. 2014. Do Employers Prefer Workers Who Attend For-Profit Colleges? Evidence from a Field Experiment.
Datta, Amit, Michael Carl Tschantz, and Anupam Datta. 2015. "Automated Experiments on Ad Privacy Settings." *Proceedings on Privacy Enhancing Technologies* 1:92-112.
Davis, Steven. 2009. "Is there a right to privacy?" *Pacific Philosophical Quarterly* 90:450-475.
de La Roncière, Charles. 1988. "Tuscan Notables on the Eve of the Renaissance." In *A History of Private Life. Revelations of the Medieval World*, edited by Georges Duby, 157-309. Cambridge: Harvard University Press.
de Waal, Frans. 2000. *Chimpanzee Politics. Power and Sex among Apes*. Baltimore: The Johns Hopkins University Press.
DeCew, Judith Wagner. 1986. "The Scope of Privacy in Law and Ethics." *Law and Philosophy* 5 (2):145-173.
Dennett, Daniel C., and Deb Roy. 2015. "Our Transparent Future." *Scientific American* 312:64-69. doi: 10.1038/scientificamerican0315-64.
Domingos, Pedro. 2015. *The Master Algorithm*. New York: Basic Books.
Dormehl, Luke. 2014. "Algorithms are great and all, but they can also ruin lives." *Wired*, 19 November 2014.
Doyle, Tony. 2009. "Privacy and perfect voyeurism." *Ethics and Information Technology* 11:181-189.
Duby, Georges. 1988a. "The Aristocratic Households of Feudal France. Communal Living." In *A History of Private Life. Revelations of the Medieval World*, edited by Georges Duby, 35-85. Cambridge: Harvard University Press.
Duby, Georges. 1988b. *A History of Private Life. Revelations of the Medieval World*. Translated by Arthur Goldhammer. Edited by Georges Duby. Vol. II. Cambridge: Harvard University Press.
Dunne, Carey. 2017. "Ten easy encryption tips for warding off hackers, the US government—and Russia. ." *Quartz*.
Dworkin, Ronald. 1997. *Taking Rights Seriously*. London: Bloomsbury Academic.
Dwoskin, Elizabeth. 2014. "FTC: Data Brokers Can Buy Your Bank Account Number for 50 Cents." 24 December 2014.

http://blogs.wsj.com/digits/2014/12/24/ftc-data-brokers-can-buy-your-bank-account-number-for-50-cents/.
Eason, Perri K., Gary A. Cobbs, and Kristin G. Trinca. 1999. "The use of landmarks to define territorial boundaries." *Animal Behaviour* 58:85-91.
Elliott, Justin, and Theodoric Meyer. 2013. "Claim on "Attacks Thwarted" by NSA Spreads Despite Lack of Evidence." *ProPublica*, 23 October 2013. https://www.propublica.org/article/claim-on-attacks-thwarted-by-nsa-spreads-despite-lack-of-evidence.
Ellis, Hattie. 2007. "The price of eggs." *The Guardian*, 18 April 2007. https://www.theguardian.com/business/2007/apr/18/supermarkets.food.
Engel, Pamela. 2013. "Police Informant Caught On Video Allegedly Framing Guy Busted For Cocaine." *Business Insider*, 26 July 2013. http://www.businessinsider.com/police-informant-allegedly-plants-cocaine-in-mans-business-2013-7.
Etzioni, Amitai. 2010. "Is Transparency the Best Disinfectant?" *The Journal of Political Philosophy* 18 (4):389-404.
Fagan, Garrett G. 2005. *Bathing in Public in the Roman World*: The University of Michigan Press.
Falchetta, Tomaso. 2016. Bulk powers equal mass surveillance equals unlawful suspicionless interference with privacy. Privacy International.
Ferenstein, Gregory. 2013. "Google's Cerf Says "Privacy May Be An Anomali". Historically, He's Right." *TechCrunch*, 20 November 2013.
Fields, Gary, and John R. Emshwiller. 2011. "Many Failed Efforts to Count Nation's Federal Criminal Laws." *The Wall Street Journal*, 23 July 2011. http://www.wsj.com/articles/SB10001424052702304319804576389601079728920.
Floridi, Luciano, and J.W. Sanders. 2004. "On the Morality of Artificial Agents." *Minds and Machine* 14:349-379.
Ford, Clellan S., and Frank A Beach. 1951. *Patterns of Sexual Behavior*. New York: Harper.
Foucault, Michel. 1977. *Discipline and Punish*. London: Penguin Books.
Fox-Brewster, Tom. 2014. "TRUSTe fined $200,000 for misleading web security seal." *The Guardian*. http://www.theguardian.com/technology/2014/nov/18/truste-fine-web-security-seals.
Freund, Paul. 1971. "Privacy: One Concept or Many?" In *Privacy: Nomos XIII*, edited by J. Roland Pennock and John W. Chapman. New York: Atherton Press.
Fried, Charles. 1970. *An Anatomy of Values*. Cambridge: Harvard University Press.
Friedersdorf, Conor. 2015. "Michael Hayden's Hollow Constitution." *The Atlantic*, 30 January 2015. http://www.theatlantic.com/politics/archive/2015/01/former-cia-and-nsa-director-nsa-doesnt-just-listen-to-bad-people/385007/.
Friedman, Lawrence M. 2007. *Guarding life's dark secrets. Legal and social controls over reputation, propriety, and privacy*. Stanford: Stanford University Press.
Gabbatt, Adam. 2013. "New York woman visited by police after researching pressure cookers online." *The Guardian*, 1 August 2013.

https://www.theguardian.com/world/2013/aug/01/new-york-police-terrorism-pressure-cooker.
Gajda, Amy. 2008. "What if Samuel D. Warren hadn't married a Senator's daughter?: Uncovering the press coverage that led to "The right to privacy"." *Michigan State Law Review* 2008 (1):35-60.
García-Guerrero, J., and A. Marco. 2012. "Sobreocupación en los Centros Penitenciarios y su impacto en la Salud." *Revista Española de Sanidad Penitenciaria* 14:106-113.
Garrett, Roland. 1974. "The Nature of Privacy." *Philosophy Today* 89:421-472.
Gavison, Ruth. 1980. "Privacy and the Limits of Law." *The Yale Law Journal* 89 (3):421-471.
Gerstein, Robert. 1978. "Intimacy and Privacy." *Ethics* 89:86-91.
Gillmor, Dan. 2014. "As we sweat government surveillance, companies like Google collect our data." *The Guardian*, 18 April 2014. http://www.theguardian.com/commentisfree/2014/apr/18/corporations-google-should-not-sell-customer-data.
Glum, Julia. 2015. "Mexico Elections 2015: Lorenzo Córdova Apologizes After Indigenous Insults from Wiretapped Conversation Go Viral." *International Business Times*. http://www.ibtimes.com/mexico-elections-2015-lorenzo-cordova-apologizes-after-indigenous-insults-wiretapped-1931122.
Goffman, Erving. 1959. *The Presentation of the Self in Everyday Life*. New York: Anchor Book.
Goodall, J. 2000. "Pride goeth before a fall." In *The Smile of a Dolphin: Remarkable Accounts of Animal Emotions*, edited by Bekoff M, 166-167. New York: Random House/Discovery Books.
Grassegger, Hannes, and Mikael Krogerus. 2017. "The Data That Turned the World Upside Down." *Motherboard*.
Greenberg, Andy. 2015. "Cops Don't Need a Crypto Backdoor to Get Into Your iPhone." *Wired*, 12 October 2015. https://www.wired.com/2015/10/cops-dont-need-encryption-backdoor-to-hack-iphones/.
Greenwald, Glen. 2013. "The crux of the NSA story in one phrase: 'collect it all'." *The Guardian*, 15 July 2013. http://www.theguardian.com/commentisfree/2013/jul/15/crux-nsa-collect-it-all.
Greenwald, Glen. 2014. *No Place to Hide*: Henry Holt and Company.
Gross, Hyman. 1971. "Privacy and Autonomy." In *Privacy: Nomos XIII*, edited by J. Roland Pennock and John W. Chapman, 169-181. New York: Atherton Press.
Gyger, M., and P. Marler. 1988. "Food calling in the domestic fowl, *Gallus gallus*: the role of external referents and deception." *Animal Behaviour* 36:358-365.
Hare, B., J. Call, B. Agnetta, and M. Tomasello. 2000. "Chimpanzwws know what conspecifics do and do not see." *Animal Behaviour* 59:771-785.
Hartshorne, Hugh, and Mark A. May. 1928. "Studies in the nature of character." In *Studies in deceit*. New York: Macmillan.
Hauser, M. 2000. "If monkeys could blush." In *The Smile of a Dolphin: Remarkable Accounts of Animal Emotions*, edited by M Bekoff, 200-201. New York: Random House/Discovery Books.

Henkin, Louis. 1974. "Privacy and Autonomy." *Columbia Law Review* 74:1410-1411.

Hicken, Melanie. 2013. "Data Brokers Selling Lists of Rape Victims, AIDS patients." *CNN*, 19 December 2013. http://money.cnn.com/2013/12/18/pf/data-broker-lists/.

Himma, Kenneth Einar. 2016. "Why Security Trumps Privacy." In *Privacy, Security and Accountability*, edited by Adam M. Moore, 145-170. London: Rowman & Littlefield.

Hirschfeld Davis, Julie. 2015. "Hacking of Government Computers Exposed 21.5 Million People." *The New York Times*, 9 July 2015. http://www.nytimes.com/2015/07/10/us/office-of-personnel-management-hackers-got-data-of-millions.html?_r=0.

Hirshleifer, Jack. 1980. "Privacy: Its Origin, Function, and Future." *The Journal of Legal Studies* 9 (4):649-664.

Hohfeld, Wesley. 1919. *Fundamental Legal Conceptions*. Edited by Walter W. Cook. New Haven: Yale University Press.

Holman, John. 2014. "Mexico's President and First Lady Face Scandal Over Lavish 'White House' Mansion." *Vice News*, 13 November 2014. https://news.vice.com/article/mexicos-president-and-first-lady-face-scandal-over-lavish-white-house-mansion.

Hopkins, Nick, Patrick Wintour, Rowena Mason, and Matthew Taylor. 2013. "Extent of spy agencies' surveillance to be investigated by parliamentary body." *The Guardian*, 17 October 2013. https://www.theguardian.com/uk-news/2013/oct/17/uk-gchq-nsa-surveillance-inquiry-snowden.

Huebert, Ronald. 1997. "Privacy: The Early Social History of a Word." *The Sewanee Review* 105 (1):21-38.

Inness, Julie C. 1992. *Privacy, Intimacy, and Isolation*. New York: Oxford University Press.

Iphofen, Ron. 2016. "Safety is more important than privacy." *Times Higher Education*, 28 April 2016.

Isikoff, Michael. 2013. "NSA program stopped no terror attacks, says White House panel member." *NBC News*, 20 December 2013. http://www.nbcnews.com/news/other/nsa-program-stopped-no-terror-attacks-says-white-house-panel-f2D11783588.

Jonas, Jeff, and Jim Harper. 2006. "Effective Counterterrorism and the Limited Role of Predictive Data Mining." *Cato Institute, Policy Analysis* 584.

Jones, W.T. 1984. "Public Roles, Private Roles, and Differential Moral Assessments of Role Performance." *Ethics* 94 (4):603-620.

Kahneman, Daniel. 2011. *Thinking Fast and Slow*. New York: Ferrar, Straus and Giroux.

Kindy, Kimberly, and Lyndsey Layton. 2009. "Integrity of Federal 'Organic' Label Questioned." *The Washington Post*, 3 July 2009. http://www.washingtonpost.com/wp-dyn/content/article/2009/07/02/AR2009070203365.html.

Klinenberg, Eric. 2012. "I want to be alone: the rise and rise of solo living." *The Guardian*. http://www.theguardian.com/lifeandstyle/2012/mar/30/the-rise-of-solo-living.

Klopfer, P.H., and D.I. Rubenstein. 1977. "The Concept of Privacy and Its Biological Basis." *Journal of Social Issues* 33 (3):52-65.
Koehler, John O. 1999. *Stasi: The Untold Story of the East German Secret Police*. Boulder: Westview Press.
Kolker, Robert. 2016. "What Happens When the Surveillance State Becomes an Affordable Gadget?" *Bloomberg Businessweek*, 10 March 2016.
Kosinski, Michal, David Stillwell, and Thore Graepel. 2013. "Private traits and attributes are predictable from digital records of human behavior." *PNAS* 110 (15):5802-5805.
Kozinski, Alex, and Misha Tseytlin. 2009. "You're (Probably) a Federal Criminal." In *In the Name of Justice*, edited by Timothy Lynch. Washington D.C.: Cato Institute.
Kramer, Matthew H. 2000. "Rights Without Trimmings." In *A Debate Over Rights: Philosophical Enquiries*, edited by Matthew H. Kramer, N.E. Simmonds and Hillel Steiner.
Kramer, Matthew H. 2010. "Refining the Interest Theory of Rights." *American Journal of Jurisprudence* 55 (31-39).
Laurent, Olivier. 2013. "Protecting the Right to Pohograph, or Not to Be Photographed." *The New York Times*, 23 April 2013. http://lens.blogs.nytimes.com/2013/04/23/paris-city-of-rights/.
Lawes, Michael J., and S. Peter Henzi. 1995. "Inter-group encounters in blue monkeys: how territorial must a territorial species be?" *Animal Behaviour* 49:240-243.
Lazar, Seth. 2015. "Risky Killing and the Ethics of War." *Ethics* 126:91-117.
Lever, Annabelle. 2012. *On Privacy*. New York: Routledge.
List, Christian. 2006. "Republican Freedom and the Rule of Law." *Politics, Philosophy & Economics* 5:201-220.
MacAskill, Ewen. 2015. "The NSA's bulk metadata collection authority just expired. What now?" *The Guardian*, 28 November 2015. https://www.theguardian.com/us-news/2015/nov/28/nsa-bulk-metadata-collection-expires-usa-freedom-act.
MacCormick, Neil. 1977. "Rights in Legislation." In *Law, Morality and Society: Essays in Honour of HLA Hart*, edited by P.M.S. Hacker and J. Raz, 189-209. Oxford: Oxford University Press.
Mackie, J.L. 1978. "Can There be a Right-BasedMoral Theory?" *Midwest Studies in Philosophy* 3:350-359.
Macnish, Kevin. 2015. "An Eye for an Eye: Proportionality and Suveillance." *Ethical Theory and Moral Practice* 18:529-548.
Macnish, Kevin. 2016. "Government Surveillance and Why Defining Privacy Matters in a Post-Snowden World." *Journal of Applied Philosophy*. doi: 10.1111/japp.12219.
Madden, Mary, Lee Rainie, Andrew Perrin, Maeve Duggan, and Dana Page. 2015. Americans' Attitudes About Privacy, Security and Surveillance. Pew Research Center.
Madden, Mary, Lee Rainie, Kathryn Zickuhr, Maeve Duggan, and Aaron Smith. 2014. Public Perceptions of Privacy and Security in the Post-Snowden Era. Pew Research Center.

Mahdawi, Arwa. 2013. "Google's autocomplete spells out our darkest thoughts." *The Guardian*. https://www.theguardian.com/commentisfree/2013/oct/22/google-autocomplete-un-women-ad-discrimination-algorithms.
Mann, Steve. 2013. *Wearable Computing and the Veillance Contract: Steve Mann at TEDxToronto*.
Mann, Steve, Jason Nolan, and Barry Wellman. 2003. "Sousveillance: Inventing and Using Wearable Computing Devices for Data Collection in Surveillance Environments." *Surveillance & Society* 1 (3):331-355.
Marmor, Andrei. 2015. "What Is the Right to Privacy?" *Philosophy and Public Affairs* 43 (1):3-26.
Mayer, Jonathan, and Patrick Mutchler. 2014. "MetaPhone: The sensitivity of telephone metadata." *Web Policy*. http://webpolicy.org/2014/03/12/metaphone-the-sensitivity-of-telephone-metadata.
McCain, Garvin, Verne C. Cox, and Paul B. Paulus. 1976. "The Relationship Between Illness Complaints and Degree of Crowding in a Prison Environment." *Environment and Behaviour* 8 (2):283-290.
McCloskey, H.J. 1980. "Privacy and the Right to Privacy." *Philosophy* 55:37-.
McDonald, Aleecia M., and Lorrie Faith Cranor. 2008. "The Cost of Reading Privacy Policies." *I/S: A Journal of Law and Policy for the Information Society* 4:543-568.
McMahan, Jeff. 2009. *Killing in War*. Oxford: Oxford University Press.
McMullan, Thomas. 2015. "The world's first hack: the telegraph and the invention of privacy." *The Guardian*, 15 July 2015. https://www.theguardian.com/technology/2015/jul/15/first-hack-telegraph-invention-privacy-gchq-nsa.
McTague, Tom. 2013. "Bullingdon Club initiation ceremony claim: New members of David Cameron's old club 'burn £50 note in front of beggar'." *Mirror*, 23 February 2013. http://www.mirror.co.uk/news/uk-news/bullingdon-club-initiation-ceremony-claim-1725912.
Merelli, Annalisa. 2015. "Charted: Terror attacks in Western Europe from the 1970s to now." *Quartz*, 25 November 2015.
Meyer, Robinson. 2015. "Could a Bank Deny Your Loan Based on Your Facebook Friends?" *The Atlantic*, 25 September 2015.
Mill, John Stuart. 1978. *On Liberty*. Indianapolis: Hackett Publishing Company.
Miller, Arthur. 1971. *The Assault on Privacy*. Cambridge: Harvard University Press.
Miller, Claire Cain. 2015. "When Algorithms Discriminate." *The New York Times*, 9 July 2015. https://www.nytimes.com/2015/07/10/upshot/when-algorithms-discriminate.html.
Mittelstadt, Brent Daniel, Patrick Allo, Mariarosaria Taddeo, Sandra Wacher, and Luciano Floridi. 2016. "The ethics of algorithms: Mapping the debate." *Big Data & Society* (July-December):1-21.
Moore, Barrington. 1984. *Privacy. Studies in Social and Cultural History*. Armonk, New York: M. E. Sharpe.
Morell, Michael. 2013. "Michael Morell: Correcting the record on the NSA recommendations." *The Washington Post*, 27 December 2013. https://www.washingtonpost.com/opinions/michael-morell-correcting-

the-record-on-the-nsa-recommendations/2013/12/27/54846538-6e45-11e3-aecc-85cb037b7236_story.html.

Mui, Ylan Q. 2011. "Little-known firms tracking data used in credit scores." *The Washington Post*, 16 July 2011. https://www.washingtonpost.com/business/economy/little-known-firms-tracking-data-used-in-credit-scores/2011/05/24/gIQAXHcWII_story.html?utm_term=.39a7ddd0ca74.

Murphy, Heather. 2013. "Why Snowden Asked Visitors in Hong Kong to Refrigerate Their Phones." *The New York Times*, 25 June 2013. https://thelede.blogs.nytimes.com/2013/06/25/why-snowdens-visitors-put-their-phones-in-the-fridge/?_r=0.

Nagel, Thomas. 1998. "Concealment and Exposure." *Philosophy and Public Affairs* 27 (1):3-30.

Newell, Bryce C. 2014. "Technopolicing, surveillance, and citizen oversight: A neorepublican theory of liberty and information control." *Government Information Quarterlly* 31:421-431.

Newman, Lily Hay. 2016. "How to block the ultrasonic signals you didn't know were tracking you." *Wired*, 03 November 2016.

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

O'Neil, Cathy. 2016. *Weapons of Math Destruction*: Penguin. Kindle Edition.

O'Neill, Onora. 2002. *A Question of Trust. The BBC Reith Lectures, 2002*. Cambridge: Cambridge University Press.

O'Neill, Onora. 2006. "Transparency and the ethics of communication." In *Transparency: The Key to Better Governance?*, edited by Christopher Hood and David Heald, 75-90. Oxford: Oxford University Press.

Olsen, Matt, Bruce Schneier, and Jonathan Zittrain. 2016. Don't Panic. Making Progress on the "Going Dark" Debate. Cambridge: The Berkman Center for Internet & Society at Harvard University.

Palazzolo, joe, and Gary Fields. 2015. "Fight Grows to Stop Expunged Criminal Records Living On in Background Checks." *Wall Street Journal*, 7 May 2015. http://www.wsj.com/articles/fight-grows-to-stop-expunged-criminal-records-living-on-in-background-checks-1430991002.

Parent, William A. 1983a. "Privacy, Morality, and the Law." *Philosophy and Public Affairs* 12 (4):269-288.

Parent, William A. 1983b. "Recent Work on the Concept of Privacy." *American Philosophical Quarterly* 20 (341-354).

Parker, Richard. 1974. "A Definition of Privacy." *Rutgers Law Review* 27.

Pasquale, Frank. 2015. *The Black Box Society. The Secret Algorithms That Control Money and Information*. Cambridge: Harvard University Press.

Peterson, Andrea. 2014. "Snowden filmmaker Laura Poitras: 'Facebook is a gift to intelligence agencies'." *The Washington Post*, 23 October 2014. https://www.washingtonpost.com/news/the-switch/wp/2014/10/23/snowden-filmmaker-laura-poitras-facebook-is-a-gift-to-intelligence-agencies/.

Pettit, Philip. 1996. "Freedom as Antipower." *Ethics* 106 (3):576-604.

Pettit, Philip. 2015. *The Robust Demands of the Good*. Oxford: Oxford Univeristy Press.
Plumer, Brad. 2013. "Eight facts about terrorism in the United States." *The Washington Post*, 16 April 2013. https://www.washingtonpost.com/news/wonk/wp/2013/04/16/eight-facts-about-terrorism-in-the-united-states/.
Poole, Steven. 2013. "Drones the size of bees—good or evil?" *The Guardian*, 14 June 2013. http://www.theguardian.com/commentisfree/2013/jun/14/drones-size-bees-good-evil.
Posner, Richard. 1981. *The Economics of Justice*. Cambridge, MA: Harvard University Press.
Powles, Julia, and Enrique Chaparro. 2015. "How Google determined our right to be forgotten." *The Guardian*, 18 February 2015. http://www.theguardian.com/technology/2015/feb/18/the-right-be-forgotten-google-search.
Precedo, José. 2015. "Espiar el móvil de la pareja: dos años y medio de cárcel." *El País*, 4 October 2015. http://politica.elpais.com/politica/2015/10/02/actualidad/1443804996_640011.html.
Quittner, Joshua. 1997. "Invasion of Privacy." *Time*, 25 August 1997.
Rachels, James. 1975. "Why Privacy is Important." *Philosophy and Public Affairs* 4 (4):323-333.
Radest, Howard B. 1979. "The Public and the Private: An American Fairy Tale." *Ethics* 89 (3):280-291.
Ramesh, Randeep. 2015. "Public bodies are releasing confidential personal data by accident, activists say." *The Guardian*, 15 July 2015. http://www.theguardian.com/technology/2015/jul/15/confidential-personal-data-release-accident-councils-nhs-police-government.
Ramos, Jorge. 2015. *Speech at TIME Gala*.
Raz, Joseph. 1984. "On the Nature of Rights." *Mind* 93 (370):194-214.
Raz, Joseph. 1988. *The Morality of Freedom*. Oxford: Oxford University Press.
Reiman, Jeffrey. 1976. "Privacy, Intimacy and Personhood." *Philosophy and Public Affairs* 6:26-44.
Ripken, Susanna Kim. 2006. "The Dangers and Drawbacks of the Disclosure Antidote: Toward a More Substantive Approach to Securities Regulation." *Baylor Law Review* 58:139-204.
Robbins, Christopher. 2011. "NYPD Uses Law from 1845 to Arrest Masked Protestors in Financial District." *Gothamist*, 19 September 2011. http://gothamist.com/2011/09/19/nypd_uses_law_from_1845_to_arrest_m.php - photo-1.
Roberts, Andrew. 2015. "A republican account of the value of privacy." *European Journal of Political Theory* 14 (3):320-344.
Rusbridger, Alan, and Ewen MacAskill. 2014. "I, spy: Edward Snowden in exile." *The Guardian*, 19 July 2014. https://www.theguardian.com/world/2014/jul/18/-sp-edward-snowden-interview-rusbridger-macaskill.

Sagar, Rahul. 2007. "On Combating the Abuse of State Secrecy." *The Journal of Political Philosophy* 15 (4):404-427.
Santos, Laurie R., Aaron G. Nissen, and Jonathan A. Ferrugia. 2006. "Rhesus monkeys, *Macaca mulatta*, know what other can and cannot hear." *Animal Behaviour* 71:1175-1181.
Savage, Charlie. 2015a. "Declassified Report Shows Doubts About Value of N.S.A.'s Warrantless Spying." *The New York Times*, 24 April 2015. http://www.nytimes.com/2015/04/25/us/politics/value-of-nsa-warrantless-spying-is-doubted-in-declassified-reports.html?_r=0.
Savage, Charlie. 2015b. *Power Wars. Inside Obama's Post-9/11 Presidency*. New York: Little, Brown and Company.
Scanlon, Thomas F. 2002. *Eros and Greek Athletics*. New York: Oxford University Press.
Schneier, Bruce. 2015a. *Data and Goliath*. London: W.W. Norton & Company.
Schneier, Bruce. 2015b. "Your TV may be watching you." *CNN*, 12 February 2015. http://edition.cnn.com/2015/02/11/opinion/schneier-samsung-tv-listening/.
Schneier, Bruce. 2016. "Security or Surveillance? (Appendix A)." In *Don't Panic. Making Progress on the "Going Dark" Debate*, edited by Matt Olsen, Bruce Schneier and Jonathan Zittrain. Cambridge: The Berkman Center for Internet & Society at Harvard University.
Sciutto, Jim. 2015. "OPM government data breach impacted 21.5 million." *CNN*, 10 July 2015. http://edition.cnn.com/2015/07/09/politics/office-of-personnel-management-data-breach-20-million/.
Segal, David. 2011. "Is Law School a Losing Game? ." *New York Times*, 8 January 2011. http://www.nytimes.com/2011/01/09/business/09law.html.
Shue, Henry. 1980. *Basic Rights*. Princeton: Princeton University Press.
Silverglate, Harvey A. 2011. *Three Felonies a Day: How the Feds Target the Innocent*. New York: Encounter Books.
Simon, Michael A. 1969. "When is a Resemblance a Family Resemblance." *Mind* 78 (311):408-4016.
Singer, Marcus G. 1959. "On Duties to Oneself." *Ethics* 69 (3):202-205.
Singer, Natasha. 2014. "Data Broker Is Charged With Selling Consumers' Financial Details to 'Fraudsters'." *The New York Times*, 23 December 2014. http://bits.blogs.nytimes.com/2014/12/23/data-broker-is-charged-with-selling-consumers-financial-details-to-fraudsters/?_r=0.
Sluga, Hans. 2006. "Family Resemblance." In *Deepening Our Understanding of Wittgenstein. Grazer Philosophische Studien 71*, edited by Michael Kober, 1-21.
Snowden, Edward. 2014. "Edward Snowden interview - the edited transcript." *The Guardian*, 18 July 2014. http://www.theguardian.com/world/2014/jul/18/-sp-edward-snowden-nsa-whistleblower-interview-transcript.
Solove, Daniel J. 2002. "Conceptualizing Privacy." *California Law Review* 90 (4):1087-1155.
Solove, Daniel J. 2006. "A Taxonomy of Privacy." *University of Pennsylvania Law Review* 154 (3):477-564.
Solove, Daniel J. 2011. *Nothing to HIde*. New Haven: Yale University Press.

Southwood, Nicholas. 2015. "Democracy as a Modally Demanding Value." *Noûs* 49 (3):504-521.
Spacks, Patricia Meyer. 2003. *Privacy: Concealing the Eighteenth-Century Self*. Chicago: The University of Chicago Press.
Stanley, Tim. 2016. "West Europe is safer now than in the 1970s. And safer than almost any other region in the world." *The Telegraph*, 25 March 2016. http://www.telegraph.co.uk/news/2016/03/25/west-europe-is-safer-now-than-in-the-1970s-and-safer-than-almost/.
Stasagage, David. 2006. "Does Transparency Make a Difference? The Example of the European Council of Ministers." In *Transparency: The Key to Better Governance?*, edited by Christopher Hood and David Heald. Oxford: Oxford University Press.
Strohm, Chris, and Del Quentin Wilber. 2014. "Pentagon Says Snowden Took Most U.S. Secrets Ever: Rogers." *Bloomberg*, 10 January 2014. http://www.bloomberg.com/news/articles/2014-01-09/pentagon-finds-snowden-took-1-7-million-files-rogers-says.
Sweeney, Latanya. 2013. "Discrimination in Online Ad Delivery." *ACM Queue* 11 (3).
Thomson, Judith Jarvis. 1975. "The Right to Privacy." *Philosophy and Public Affairs* 4 (4):295-314.
Traynor, Ian. 2015. "EU referendum: Brussels tells evasive Cameron to spell out agenda for talks." *The Guardian*.
Veyne, Paul. 1987. "The Roman Empire." In *A History of Private Life. From Pagan Rome to Byzantium*, edited by Paul Veyne. Cambridge: Harvard University Press.
Waddell, Kaveh. 2016. "How Algorithms Can Bring Down Minorities' Credit Scores." *The Atlantic*, 2 December 2016.
Waldron, Jeremy. 1989. "Rights in Conflict." *Ethics* 99 (3):503-519.
Waldron, Jeremy. 2003. "Security and LIberty: The Image of Balance." *The Journal of Political Philosophy* 11 (2):191-210.
Wang, Peggy. 2009. "No Eye-Contact Glasses." *BuzzFeed*. http://www.buzzfeed.com/peggy/no-eye-contact-glasses/ - .aaxp96d5m.
Warren, Samuel D., and Louis D. Brandeis. 1890. "The Right to Privacy." *Harvard Law Review* IV (5):193-220.
Wasserstrom, Richard. 1978. "Privacy: Some Arguments and Assumptions." In *Philosophical Law*, edited by Richard Bronaugh, 148-166. Westport, CT: Greenwood Press.
Watson, Gary. 2013. Moral Agency. In *The International Encyclopedia of Ethics*.
Webb, Diana. 2007. *Privacy and Solitude in the Middle Ages*. London: Hambledon Continuum.
Welsh, Brandon C., and David Farrington. 2004. "Surveillance for Crime Prevention in Public Space: Results and Policy Changes." *Criminology and Public Policy* 3 (3).
West, Lindy. 2016. "The macabre truth of gun control in the US is that toddlers kill more people than terrorists do." *The Guardian*, 13 March 2016. https://www.theguardian.com/commentisfree/2016/mar/13/the-macabre-truth-of-gun-control-in-the-us-is-that-toddlers-kill-more-people-than-terrorists-do.

Westin, Alan F. 1970. *Privacy and Freedom*. London: Bodley Head.
Wheelwright, Geof. 2016. "What are the big tech companies lobbying for this election?" *The Guardian*, 26 September 2016. https://www.theguardian.com/technology/2016/sep/26/tech-news-lobby-election-taxes-tpp-national-security.
Whiten, A., and R.W. Byrne. 1988. "Tactical deception in primates." *Behavioral and Brain Sciences* 11:233-273.
Wright, Lawrence. 2008. "The Spymaster." *The New Yorker*.
Yadron, Danny. 2016. "'Worth it': FBI admits it paid $1.3 to hack into San Bernardino iPhone." *The Guardian*. https://www.theguardian.com/technology/2016/apr/21/fbi-apple-iphone-hack-san-bernardino-price-paid.
Zahavi, Amotz, and Avishag Zahavi. 1997. *The Handicap Principle. A Missing Piece of Darwin's Puzzle*. Oxford: Oxford University Press.
Zenko, Micah. 2012. "Americans Are as Likely to Be Killed by Their Own Furniture as by Terrorism." *The Atlantic*, 6 June 2012.
Zuboff, Shoshana. 2015. "Big Other: Surveillance Capitalism and the Prospects of an Information Civilization." *Journal of Information Technology* 30:75-89.