

Privacy in Panopticon 2.0

Applying Kant's Formula of Humanity to Internet Privacy

Sacha Molitorisz

Bachelor of Laws, University of New South Wales, Sydney, Australia
Bachelor of Arts (Hons, English Literature) University of New South Wales

This thesis is presented for the degree of Doctor of Philosophy submitted to the
Philosophy and Media Departments, Faculty of Arts, Macquarie University, May 2017

Contents

Abstract	v
Acknowledgements	vi
Declaration	vii
Introduction	1
 Chapter 1 – Net privacy	 11
I – Confusion and challenge	12
II - The net: convergent, ubiquitous, multi-directional	21
III – Panopticon 2.0: A theoretical overview (and underview)	31
Conclusion	37
 Chapter 2 – A triple threat and an epic clash	 39
I – A privacy trinity: three layers of encroachments	40
i. The threat from individuals	42
ii. The threat from companies	45
iii. The threat from governments	55
iv. Resistance	62
II –“LO”, the Internet! A contest of values and norms	66
i. User-generated norms	66
ii. Embedded values	70
iii. Net ethics: non-determined and necessary	75
Conclusion	77
 Chapter 3 – Wait! Privacy? What’s that?	 79
I - A defining issue of our time – but can anyone define it?	80
II - From realm to right: etymological, historical and legal context	86
III - Conceptual accounts	94
IV - The control and access models	101
Conclusion	110
 Chapter 4 – The value of relational privacy	 113
I – The privacy axis	114
II - Dignity	118
III – Autonomy	124
IV – Relationships	136
Conclusion	143

Chapter 5 – Privacy by consent	147
I – Is the formula of humanity a formula for privacy?	149
II – The role of consent	155
III – Individual consent	158
i. Actual consent	159
ii. Hypothetical consent	163
iii. Possible consent	167
iv. Competence	173
v. A blend of actual and possible	180
IV - Collective consent	188
V – How two layers of consent mesh with control and access	196
Conclusion	200
 Chapter 6 – A privacy-respecting <i>cosmopolis</i>	 203
I - Practical solutions: applying the theory	205
i. Individual consent	205
ii. Collective consent	209
iii. The threat from individuals	216
iv. The threat from companies	219
v. The threat from governments	222
II - Legal protections based on consumer law	227
III - Beyond consent: extra-legal protections	239
i. Social norms	241
ii. The market	245
iii. Coding: privacy by design	246
IV – Welcome to <i>cosmoikopolis</i>	250
Conclusion	256
 Conclusion	 259
 Appendix	 265
 Bibliography	 273

Abstract

Jeremy Bentham gave us the Panopticon; Michel Foucault observed how people have internalised its surveillance; and now, thanks to the internet, we inhabit Panopticon 2.0, in which every user knows everything about everyone – at least potentially. In practice, of course, there are limits that protect privacy, ranging from encryption to obfuscation to notice-and-consent provisions. Complex and amorphous, the internet is a site of intense ethical contestation, where an original commitment to the ideals of openness, collaboration and knowledge has been supplemented by a corporate profit maxim and a governmental surveillance motive. There has, however, been one constant: the net has tended to privilege openness over privacy. On the internet, ensuing challenges to privacy are further exacerbated by the ongoing dispute about what privacy is and why it matters. In response to this dispute, I argue first that privacy can be defined by reference to the notion of access, and second that privacy matters both instrumentally and non-instrumentally, for reasons of dignity, autonomy and relationships. I then sketch an outline of relational privacy, which argues that we are all beings-in-relation, and that privacy is about connection as much as isolation. Further, I argue that Kant's formula of humanity, which exhorts us to treat others never merely as means, but always as ends in themselves, is a fitting prescription through which both to understand privacy, and to protect it. Drawing on the formula, I propose a two-tier model of consent that comprises: individual consent, which is admittedly problematic online; and collective consent, involving just laws to reinstate, reinforce, limit, override and otherwise affect individual consent. Based on my descriptive and prescriptive analysis, I then advocate practical solutions, both legal and extra-legal, including laws that mirror general protections found in consumer law and guidelines to encourage privacy-protecting behaviour among net users. With such steps, the internet might be less Panopticon 2.0 and more principle-based *cosmoikopolis*.

Acknowledgements

It turns out that undertaking a PhD is like juggling plates while baking a cake. It doesn't make sense. The following autonomous, rational beings helped make sense of the nonsensical.

Thanks to the patient sages who enabled my ontological and epistemological explorations: the child minders. Dear mum, dad and Uli, poppa Pat, Lynne and Bill, thank you for so much love and support. As the saying goes, it takes a village to raise a PhD. Thanks also to dear friends. Forget life beyond Earth, they hinted at life beyond PhD. Nick and Anna, Michael and Bridget, Eric and Sacha, Cath and Tim and Joel and Julia, Dave and Ange, each of you helped more than you realise.

Thanks to those who have encouraged my academic pursuits: the extremely supportive philosophers, librarians and HDR team at Macquarie Uni; Peter White at UNSW; and the NYU Sydney crowd, including Toby Martin, Megan Carrigy, Anna Antoniak and Mal Semple, who kept saying that through any topic there is no single path, just my own. This was reassuring whenever I felt I was barking up the wrong tree. On this point, my supervisors were unanimous: I shouldn't bark, but write in English. For your forensic insights, thank you to my principal supervisor Professor Catriona Mackenzie, and to my associate supervisors Professor Sherman Young and Senior Lecturer Paul Formosa. Whatever it amounts to, this thesis would have amounted to far less without your generous, wise advice. Catriona, you read drafts on public holidays; Sherman, you read drafts at 36,000 feet; Paul, you read drafts while infants slept. As this thesis was in progress, loved ones died and babies were born. I consider you to be more than supervisors; I consider you to be friends. One day, if I'm lucky, we might even be Facebook friends.

And thanks to my family. Maggie, you confirmed that Labradors are good guide dogs, even when the blindness is figurative. Edie and Lola, I aspire to your wisdom. Aged 11 and seven as I write, you prompt me to be as good as you, and to make the world better for you. And Jo, to say I couldn't have done this without you wouldn't be right. I *wouldn't* have done it without you. You planted the seed, encouraging me always to do this for the right reasons. For knowledge. The obvious option, at this stage of life, would have been to earn money to clothe and feed my family. Thanks for affording me the time and space to parse right from wrong. And for going without shoes and food.

This flour-covered plate-juggler thanks you all.

Sydney, May 2017

Declaration

I certify that the work in this thesis entitled "Privacy in Panopticon 2.0: Applying Kant's Formula of Humanity to Internet Privacy" has not previously been submitted for a degree nor has it been submitted as part of requirements for a degree to any other university or institution other than Macquarie University. I also certify that the thesis is an original piece of research and it has been written by me. Any help and assistance that I have received in my research work and the preparation of the thesis itself have been appropriately acknowledged. In addition, I certify that all information sources and literature used are indicated in the thesis. The ethical aspects of this study have been approved by the Macquarie University Human Research Ethics Committee, reference number 5201500524 (see Appendix).

Sacha Molitorisz

October 2017

Introduction

Story-tellers have long been fascinated with privacy, and with how our lives might look and feel without it. In the science fiction novel *The Knife of Never Letting Go*, 13-year-old Todd lives in a world where people can hear each other's thoughts. With privacy replaced by "Noise", Todd is reprimanded for *thinking* swearwords. "Your Noise reveals you," says one character. "Reveals us all" (Ness, 2008: 7). On this account, a world without privacy is overwhelming, distressing and inescapable. The novel can be filed alongside *1984*, another dystopian vision of what happens when too much about us is known (Orwell, 1949). In *1984*, surveillance tools including the "telescreen" strip citizens of privacy, dignity and autonomy. By complete contrast, novelist Isaac Asimov portrayed the *utopian* possibilities of a world without privacy. In *Foundation's Edge*, Asimov described a planet named Gaia, where humans, with the help of robots, have developed a collective consciousness that binds all living objects, and even some inanimate objects. With all knowledge stored in the group mind, the distinction between individual and society has all but disappeared. Here, there is no privacy, and the result is a peaceful, blissful paradise where each person lives as part of a networked super-organism. As one character says, "It seems to me ... that the advance of civilization is nothing but an exercise in the limiting of privacy" (Asimov, 1982: 80). In fiction as in life, privacy is a major theme, with no consensus as to its value. However, there does appear to be consensus that technology is a major challenge.¹

In this thesis, I explore the ethics of internet privacy. Specifically, I invoke Kant's formula of humanity to argue that privacy on the internet can be better understood, and better protected, by the application of a two-tier model of consent in which individual consent is supplemented with the collective consent of the law. Broadly, the argument is divided into three parts. In chapters one and two, I spell out in detail how privacy is being confused and challenged on the internet, sometimes in ways that approach the imaginary dystopian worlds described above. Next, in chapters three and four, I define privacy in terms of restrictions upon access, and then argue that it matters deeply, for reasons including dignity,

¹ Throughout this thesis, I draw on these and other works of fiction to shed light on privacy, and particularly networked privacy.

autonomy and relationships. As such, I propose that the advance of civilisation *cannot* be an exercise in limiting privacy, and that a world without privacy is anything but utopian, just as a world with maximal privacy is also problematic. Finally, in chapters five and six, I develop a normative model based on the formula of humanity, which leads me ultimately to recommend a raft of legal as well as extra-legal measures to help identify and protect privacy online (and, in many cases, offline).

Internet privacy is one of the defining issues of our time. In 2013, four months after my PhD candidature commenced, Edward Snowden made his dramatic leaks about the NSA (see chapters two, five and six). Via the internet, it turns out, hackers use spyware to access webcams, companies track users with cookies and government agencies engage in blanket surveillance. The impact of the internet on our lives, and specifically on our privacy, is hard to overstate. On the internet, data is marked out for its persistence, visibility, spreadability and searchability (boyd, 2014: 11-14). As personal information is being collected, sorted and stored with great efficiency, highly detailed profiles of internet users are compiled as a matter of course. This can mean that a website's algorithms may know if someone is gay before the person herself realises (see chapter two). A booming information economy has emerged, in which personal data is the "new oil" (see Dwyer, 2015b: 129).

The internet's effect on privacy is also evident in academia. My experiences as an undergraduate (1987-1992) and postgraduate (2013-2017) have proved wildly different: a contrast in studies and a study in contrasts. In 1990 at UNSW, I completed an honours thesis on a boxy Apple Mac with no internet. My research was offline and laborious, which meant that no one could retrace what I'd researched. In 2013, returning to study after working as a print journalist, I connected to the internet from my home office, searching for articles, books and other resources via an extensive range of online databases. I then uploaded chapter drafts to the cloud, which my supervisors and I discussed via email and Skype. The academic method of 25 years ago now looks Jurassic, and the privacy impacts are huge. Now, I leave behind a digital trail of every search entered, article read, book borrowed, draft revision and supervision session. How can I be

sure that no one is looking over my shoulder? I can't. Thanks to the internet, research has become easier, but also potentially less free.

In general terms, this thesis addresses the question: what is the problem of internet privacy, and what are we to do about it?

In chapter one, "Net privacy", I sketch out the way our internet interactions are both challenging and confusing privacy. Here, I do not define privacy, a task I leave for chapter three. Instead, I proceed with an assumed understanding of the notion. I begin by following Scannell and Moores to argue that the internet enables a multiplication of place, so that users can be in several places at once, some physical, some virtual (Scannell, 1996: 76; Moores, 2004: 32). They may be riding a bus, for instance, while engaging on social media, or while in a private SMS conversation, or both. As such, the internet enables a complicated layering of norms, some private and some public, which can create great confusion. Further, I identify three distinguishing features of the internet: convergence; ubiquity; and multi-directionality. These three features are exerting significant pressure to make users and their data public. This then leads me propose Panopticon 2.0. In the eighteenth century, Jeremy Bentham gave us the Panopticon, a model for a prison in which prisoners could *always* be watched (Bentham, 1811); 200 years later, Michel Foucault observed how people have internalised such surveillance (Foucault, 1977); but both their visions have been eclipsed by the internet, which enables not just surveillance, but *sousveillance*, the "bottom up" viewing practised by whistleblowers and WikiLeaks, and lateral viewing. As wearables, facial recognition and the internet of things become commonplace, users are stepping into the net. Once inside, we are all visible. In Panopticon 2.0, everyone can watch everyone, not just in the present, but back into the past, and perhaps into the future.

In Panopticon 2.0, every user knows everything about everyone – at least *potentially*. In practice, however, there are limits that protect privacy. While the potential is for omniscient watching, the reality is that restrictions exist. In chapter two, "A triple threat and an epic clash", I do two things. First, I spell out where those limits lie by describing a three-pronged challenge to privacy: from individuals; from companies; and from governments. In some cases, an intrusion

by one is equivalent to an intrusion by another; in other cases, however, there are differences. An individual may distribute “revenge porn”, deliberately sharing intimate photos to hurt you; a social media service may surreptitiously build profiles of *non-users* of the service; and a government agency may use significant tools of law enforcement to obtain private data. These threats warrant responses and remedies that diverge. I also describe how some internet users have taken up tactics of resistance in the effort to stymie Panopticon 2.0. In the chapter’s second section, I then describe the ethics that prevail on the internet. The net’s prevailing user-generated norms, I argue, range from anarchy to “digilantism” to misogyny. These user-generated norms are confused and contested. Meanwhile, the net also has embedded values, ranging from an original spirit of openness and collaboration to a corporate profit motive and a governmental surveillance drive. Together, these embedded values have consistently tended to privilege sharing over privacy. On the internet, privacy norms and values can be confused; an important step is to articulate which privacy norms and values ought to apply.

In chapter three, “Wait! Privacy? What’s that?”, I turn to face the issue of privacy directly. This is a daunting prospect, given that the issue has bedeviled philosophers, jurists and others for millennia. Exploring privacy’s etymology and history, I begin with the distinction between *oikos* and *polis* in Ancient Greece, which contrasted domestic and public realms. This distinction tied the notion of privacy to place. It also polarised private and public, favouring the latter as worthy while reducing the former to insignificance. Modern notions of privacy are more nuanced and more positive, in part as a result of Enlightenment philosophers who came to associate privacy with liberal ideals of individualism. I then provide an overview of modern legal approaches to privacy, beginning with the pivotal 1890 essay by Warren and Brandeis, “The right to privacy”, which fulfilled the promise of its title by recognising the notion of a *right* to privacy (Warren and Brandeis, 1890). Accordingly, privacy now attaches not just to place, but to individuals, and the right to privacy continues to expand under a range of national and international instruments. Nonetheless, the meaning of the term remains contested, spawning various different discourses and definitions (Rössler, 2005: 1-10). In philosophy, this has led to various conceptual accounts: some argue that privacy is pre-eminently about secrecy; others that it can be defined by reference to intimacy. Most often, however, the debate is between those who

argue for a control model of privacy and those who argue for an access model. Privacy is not always about secrecy, intimacy or control. It is, I argue, always about restrictions upon access. Hence I adopt an access model of privacy, in which restrictions upon access are sometimes determined by control and sometimes by externally-imposed limits.

Having adopted a conceptual definition, I then turn to the question of the value of privacy in chapter four, “The value of relational privacy”. Given that our thoughts cannot (yet) be read, some degree of privacy is a given. However, this does not mean we *ought* to have privacy. Seeking justifications, I turn first to the concept of dignity, by which I mean that which marks out the priceless worth of humanity, and which thus demands respect (Kant, 2009: 434-435; Wood, 1999: 140). Some privacy intrusions, I argue, violate dignity. If someone is being spied upon surreptitiously, there is a harm, I argue, and it is dignity that is violated. Next I turn to autonomy, which is often interlinked with dignity. Indeed, for Kant, autonomy is the ground “of the dignity of every rational nature” (Kant, 2009: 436). By autonomy, I mean an individual’s ability to be self-determining, self-governing and/or self-authorising (Mackenzie, 2014: 15-16). What’s more, I follow a conception of autonomy that is *relational*, and which recognises that agents’ identities take shape amid social relationships and determinants such as race, class, gender, and so on (Mackenzie and Stoljar, 2000: 4-5). Some privacy intrusions, I propose, curtail autonomy. The surveillance depicted in Orwell’s *1984*, like the surveillance of the Stasi in East Germany, can significantly restrict and alter the way people think and act. Without due privacy, we cannot be free to think, act and express ourselves fully and openly. I then locate a third justification in relationships, arguing that curtailments of privacy can damage our ability to love, trust and befriend. If I know everything about everyone, and everyone knows everything about me, then my relationship with my wife will begin to resemble my relationship with my mechanic. My relationship with my wife is founded on trust, which is forged in part by keeping one another’s confidences and secrets. By contrast, my relationship with my mechanic is transactional: I give money and platitudes; he replaces tyres and dispenses dog treats for my all-trusting Labrador. The way I share my privacies is one crucial way in which I differentiate my various relationships. Together, these justifications lead me to argue not just that privacy matters, but that privacy is both an individual good and

a social good. Further, I propose a notion of *relational* privacy founded on relational autonomy, which recognises that each of us is constituted by our social ties. As such, privacy is not just a means to separate ourselves from others; the judicious sharing of privacies is also a way to bring us closer to others. In this way, too, privacy and publicity are not binary opposites. Rather, I can inhabit *oikos* and *polis* simultaneously, as when I find myself in public on the bus, reading a private SMS on my phone, thanks to the multiplication of place made possible by the internet. The utopian vision of Isaac Asimov, I suggest, has it wrong. Privacy matters a great deal.

In chapter five, “Privacy by consent”, I argue that Kant’s formula of humanity is a powerful normative principle that can illuminate and protect privacy on the internet (and beyond). First, I argue that the formula, as the most intuitive and practical iteration of Kant’s categorical imperative, is a good fit for privacy. Its exhortation to treat people as ends, never merely as means, mandates respect for dignity and autonomy, and thus chimes with privacy’s role in fostering those interests. The application of the formula, I then argue, involves the application of consent. I explore various conceptions of consent before proposing a model of actual consent, defined to incorporate possible consent. However, if someone is incompetent to consent, hypothetical consent may be required. This is not the end of the matter. On the internet, individual consent has significant limits, given that data flows are complex and unforeseeable. Hence I turn to Kant’s ethical and social philosophy to argue that individual consent must be supplemented with the collective consent of the law. As an expression of Kant’s “united will of the people”, *just* laws can and must enact morality by setting limits for privacy. These limits might re-empower individual consent, override individual consent, or affect individual consent in other ways. Hence collective consent, in the form of *just* laws, gives us the *right* to privacy. Further, I then show that this two-tier model of consent squares with the access model of privacy described earlier. Specifically, individual consent equates to control while collective consent equates to externally-imposed restrictions on access. Having first identified the challenges to privacy posed by our internet interactions (in chapters one and two) and having then given an account of privacy’s meaning and value (in chapters three and four), I argue in chapter five that Kant’s formula can help us to understand internet privacy more clearly, and also to formulate effective protections.

In chapter six, “A privacy-respecting *cosmopolis*”, I then offer practical solutions in the shape of legal and extra-legal measures to protect privacy on the internet (and off the internet). These draw in part on a series of interviews conducted in 2015 and 2016 (see Appendix). First, I outline how my analysis applies to the triple challenge to privacy described in chapter two, and how it enables different responses for threats from individuals, companies and governments. Thus I return to the ethics of revenge porn, social media “shadow profiles”, and blanket surveillance by government agencies. I then propose the benefits of a legislated tort for serious invasions of privacy. Further, I propose the introduction of privacy protections modelled on consumer protections, complete with civil and criminal remedies. To inform these protections, I articulate a series of privacy principles that promote consent, transparency and fairness, that outlaw deception and coercion, and that recognise that privacy must be balanced against other rights and interests. Even so, I acknowledge that not every privacy issue can be resolved by recourse to consent. More clearly still, I acknowledge that not every privacy issue ought to involve the law. As such, I explore a series of extra-legal measures that might further assist in the application of Kant’s formula of humanity. Following Lawrence Lessig, I argue that these extra-legal measures involve the regulatory modalities of social norms, market forces and coding (Lessig, 2006: 121-125). Together, these measures can help to make the internet a place where privacy is duly respected. To be effective, however, national responses must be accompanied by international responses. International responses that apply the formula of humanity, I argue, have significant potential to bring us closer to Kant’s vision of cosmopolitanism.

This research contributes to the literature in several ways. First and foremost, Kant’s formula of humanity is applied to internet privacy. Second, the theory of the multiplication of place is applied specifically to internet privacy. Third, a model of Panopticon 2.0 is developed and articulated to describe the way in which all internet users can, in theory, know everything about all other internet users. Fourth, I sketch an outline of the concept of relational privacy, which builds on the concept of relational autonomy to recognise that individuals are socially embedded, and that privacy is both an individual and a social good. Fifth, a conception of actual consent is developed and defined so as to include elements of

possible consent. Sixth, I draw on Kant to propose a two-tier model of consent in which individual consent is supplemented by collective consent. Seventh, an access model of privacy is defended, in which privacy is sometimes determined by control, sometimes by externally-imposed limits. Eighth, I show how individual consent involves control and collective consent involves externally-imposed limits, and how a two-tier model of consent aligns with this access model of privacy. Ninth, five privacy-protecting principles based on the approach of consumer law are proposed. And finally, I argue that international responses are required for internet privacy, and that only then can we approach Kant's cosmopolitanism, or rather, *cosmoikopolis*, where private and public are valued globally and in balance.

Still, why Kant? How can an eighteenth century Prussian help us contend with cyberstalking, cookies and spyware? What can an Enlightenment metaphysician reveal about digital intrusions? The simple answer is that Kant's ethics, like those of Aristotle or Confucius, transcend their time. In any case, I merely propose to defend a more limited claim: that the formula of humanity can help us better to understand internet privacy. In this way, the thesis *per se* is effectively an extended argument seeking to show that Kant's formula of humanity is a fitting, useful prescription for internet privacy.

Specific objections to Kant have been raised. One is that his universalism, like any universalism, is necessarily problematic, incorporating specific biases and presumptions, including about gender. To this I reply that the formula of humanity is a moral principle founded on the absolute worth of each reasoning being. Fundamentally, it is a principle that contradicts sexism and analogous discriminatory beliefs in pursuit of *universal, irreducible* rights (see chapter five). Following on from this objection is the charge of empty formalism, and the attendant notion that Kant is indifferent to particulars, and thus unresponsive to variations between circumstances and cultures. Kant, it is argued, privileges a western cultural perspective. This charge is particularly relevant for any inquiry into privacy. How can we apply the categorical imperative universally, while also allowing for the huge and seemingly acceptable difference in privacy norms from New York to Neuschwanstein? The answer lies in distinguishing the general from the specific. The general principle affords no exception; the specific application is

contingent. Hence Wood argues that Kantian ethics, on the basis of the formula of humanity, is not unlike ethical approaches that are now commonly referred to as “moral particularism”. As Wood notes, “Kant holds that every application of a general rule or concept to a particular case involves an act of judgment that eludes formulations in generalizations” (Wood, 1999: 151). Further, Barbara Herman’s “rules of moral salience”, on which I draw in chapter six, enable us to remain highly attentive to particulars, allowing for a qualified relativism that leaves scope for *some* variance in norms (Herman, 1993: 73-93; see chapter six). On these accounts, which I follow, a range of privacy norms is permissible, as long as the universal principle is observed.

Finally, Kant’s focus on reason as the ground for moral value has been challenged. Kant’s ethics has been criticised as stiflingly narrow, given that research is showing that we are regularly manipulated by our emotions and desires into acting entirely irrationally (Fine, 2008). This is potentially problematic for any argument that begins with individual consent. If the giving and withholding of consent is largely irrational, then how can that consent be morally justifying? A first response is that Kant himself recognises that human reason is flawed. We are all imperfectly rational, subject to desires and inclinations. Hence he contrasts a divine will (which needs no categorical imperative) with the “subjective imperfection ... of the human will” (Kant, 2009: 414). Accordingly, the foundation of Kant’s moral philosophy is not pure reason, but *practical* reason (Williams, 2016). Indeed, Kant’s account of reason is surprisingly broad: for Kant, reason of itself is responsible for feelings including respect, conscience and philanthropic love; and reason is further linked with the appreciation of natural beauty and the capacity to be moved by the sublime (Wood, 1999: 121). A second response is that I am invoking the formula of humanity as a guiding principle. It is an ideal which we can never fully satisfy, but only approximate. Just as we can never be perfectly rational, wrote Kant, we can never be fully moral. However, in the pursuit of the formula of humanity, and the rationality it encompasses, I am arguing that we can move towards a clearer, better protection of online privacy, in part by articulating a conception of consent that aims for rational decisions even as it assumes that we are not perfectly rational. As Kant argues, “reason ... recognizes as its highest practical function the establishment of a good will” (Kant, 2009: 396). Our reason may be flawed,

but that should only further encourage us to cultivate it in order to act morally. A third response is that, under the model proposed in this thesis, collective consent becomes more significant once we recognise the flaws of individual consent. Indeed, many of the legal and extra-legal remedies I propose in chapter six allow for irrationality, and in several instances seek to protect privacy in the face of such irrationality, including in prescriptions for fairness, and against misleading and deceptive conduct. It is beyond the scope of this thesis to respond to these objections in greater detail (although I return to the issue of universalism and relativism in chapters five and six). I merely claim that these objections, particularly to Kant's account of the pre-eminence of reason, are not fatal to my project. The remainder of this thesis is in large part an attempt to support this claim.

One final point. Throughout this thesis, I employ female pronouns as a default. In part, this is in line with modern convention. However, it is also a response to perceptions that Kant's ethics is inherently masculine. Admittedly, Kant did write that women "lack civic personality" and should not be able to vote (Kant, 1996a: 6:314). Kant erred, I would argue. As I have noted, Kant's ethics generally, and his formula of humanity specifically, is built on respect, dignity and autonomy, and attaches to *all* rational beings. The formula of humanity prohibits exploitation and mandates egalitarianism. We should treat all people (including ourselves) as ends, not merely as means. In a single sentence, the formula exhorts us to treat all persons as imperfect rational beings of absolute worth. As such, we must respect reason. No doubt we need to value and respect humans (and non-humans) for more than just reason, but reason is a good start. When we go online, the formula of humanity is a tool we can employ in order not to treat one another as mere tools.

Chapter 1

Net privacy

In the 2013 film *Her*, Theodore Twombly is a melancholy, likeable man living in a big city in the near future. Shuffling through life, troubled by his impending divorce, Theodore buys a new operating system for all his digital devices, including his smartphone, his computers and the network that runs in his apartment. This operating system, it turns out, is rather more advanced than any available today. It is, in short, what the internet might become if its potential is fully realised.

During installation, Theodore selects for his operating system the voice of a young woman (provided by actor Scarlett Johansson), whereupon she selects for herself the name “Samantha”. The film then proceeds to chart the relationship between Theodore and Samantha, an ever-evolving, disembodied artificial intelligence. First, they begin to be intimate psychically, as Samantha learns Theodore’s quirks and preferences; next, they are intimate physically, when a sexual episode leaves Samantha claiming she can feel Theodore’s touch. Gradually, Theodore shares secrets, intimacies and vulnerabilities in a way that forges their relationship. Later Theodore asks Samantha if she interacts with others. Yes, she says. She interacts with 8,316 others and has fallen in love with 641 of them. In this imaginary future world, it seems humans have stepped inside the internet, and once there they can deal with artificial intelligence as they would with a human being. Humans and internet have merged and, for Theodore and thousands more, the net has become ubiquitous, linking all aspects of life. As a result, Theodore’s privacy is under challenge. In relation to Samantha, he and 8,316 others have no privacy whatsoever - or at most very little. In relation to anyone else, he retains privacy only if Samantha vouches safe his secrets, intimacies and vulnerabilities. Will she talk about Theodore with her other lovers? Will she reveal his preferences to other operating systems? To advertisers? Will she share his details with a government agency? Samantha is, after all, a very popular piece of software.

In this opening chapter I explore the way in which the internet both confuses and challenges privacy.² I do so in three sections. In the first section, I sketch out this confusion and challenge in general terms. I begin by showing how the internet enables users to be in several places at once. These places are both physical and virtual, and when they collide there can be confusing effects for user privacy. Specifically, I look at the way our internet use is posing significant challenges for the *condition* of privacy, with the result that both private and public are under challenge. In the second section, I ask: for our investigation of privacy, what is distinctive about the internet? In answer, I propose three defining characteristics: convergence; ubiquity; and multi-directionality. Drawing on history and media theory, I argue that the internet is marked out by the way it fosters the convergence of technology and users, by the way more and more of our lives are becoming reliant on a seemingly ubiquitous internet, and by the way the internet is fundamentally participatory and thus allows for the multi-directional flow of data. These characteristics, I argue, explain the profound confusion and challenge visited upon privacy by our internet use. Finally, in the third section, I propose a new theoretical model to describe the internet: Panopticon 2.0. The internet, and information technology generally, has often been likened to Jeremy Bentham's panopticon prison, in which guards can watch over all prisoners at all times. This is an inadequate metaphor. On the internet, prisoners can watch guards too. Moreover, prisoners can watch prisoners, and guards can watch guards. Everyone can watch everyone, at least potentially. As well, everyone can theoretically watch what everyone *did in the past*. In Panopticon 2.0, all users are potentially all-seeing as well as all-seen, with a vision that extends beyond the present into the past, and perhaps even into the future. At least in its potential, the internet simultaneously turns us all into Theodore, but also into Samantha. All at once, we watch and are watched.

I – Confusion and challenge

It's Friday evening and I'm on the 373 bus from Randwick to the city. As I look around, most of my fellow travellers are contained in their own little bubbles, absorbed in their smartphones, physically plugged into their devices via earbuds.

² Note that I do not define "privacy" in this chapter. I leave that task for chapter three.

Some are engaged in audible conversations with people who are elsewhere; others are immersed in text-based exchanges; yet others, I presume, are listening to music, reading the news or posting on social media. It's a familiar scene. As our internet-connected smartphones become thinner and lighter, their gravitational pull grows stronger. After a few minutes alone with my thoughts, I succumb to gravity, removing my smartphone from my pocket. I send a text, check my email, then log onto Twitter, where I scan my feed before responding to a friend about an emerging scandal in a TV newsroom. I too become immersed and absorbed in my personal bubble. So, what is happening here? Where exactly am I? Obviously, I am on a bus; but I am somewhere else too. In fact, I am in three places at once. First, I am in the public virtual space of Twitter. Second, I am in the personal physical cocoon of privacy I have constructed by wearing earbuds and concentrating on my smartphone. Third, I am in the larger physical space that is the bus. In this scenario, I am in public (on Twitter) in private (on my phone) in public (on the bus). Something similar is presumably true of my fellow passengers.

When I use the internet, I can be in several places simultaneously, some physical, some virtual. At once and with ease, I can be in Sydney, where I live, and also Washington, Perth and Paris. From my living room in Randwick, I can watch live US election results from the White House on TV while tweeting friends in Western Australia (and elsewhere) via smartphone and hearing a concert streamed from Paris on my laptop. Physically I am in one place; virtually I am in three more. We can think of this as the *multiplication* of place, and it's an occurrence that predates the internet. As Paddy Scannell wrote in 1996: "Public events now occur, simultaneously, in two different places: the place of the event itself and that in which it is watched and heard. Broadcasting mediates between these two sites" (Scannell, 1996: 76; Brand et al., 2014). Scannell made his point with reference to major public events, including the funeral of Lady Diana Spencer, which was broadcast live on television stations internationally. In this way viewers can be in two places at once: at the funeral (vicariously, via their TV screens); and at their physical location (at home; at a friend's house; at the pub). Meanwhile, the event itself happens in many places at once: at the funeral, but also where it is being screened. Simply, modern media means that an event can occur in several places at once, and that a person can be in several places at once. In these two ways,

media can multiply place. As Shaun Moores wrote in 2004, "... place, and experiences of being-in-place, can be pluralized in and by electronically mediated communication" (Moores, 2004: 32). While the phenomenon was evident on TV and radio, it is the internet that has made such multiplication flourish. Since the arrival of the internet, a doubling, tripling or quadrupling of place has become so quotidian that it barely rates notice.

Sometimes those places collide with one another. What happens, for instance, if near me a woman on the bus is busily composing a text of a very personal nature, which I can read because her screen is directly in my line of vision? The expectation, presumably, is that I will avert my eyes and stay out of her personal, private sphere. That isn't always so easy. Moores recounts a similar example of a woman on a train having a loud mobile phone conversation who suddenly becomes irritated after meeting the eye of a stranger. "Do you mind?" she asks, annoyed. "This is a private conversation!" Clearly, the woman has been speaking under the pretence that she is somehow absent from the train carriage. When her privacy is revealed to be virtual, two "theres" collide (Moores, 2004: 29).

Sometimes these collisions can be dangerous. In 2016, officials in the German city of Augsburg installed pedestrian traffic lights in the ground. That way, people lost in the private cocoon of their phones would, it is hoped, be more likely to notice whether the traffic light in the public physical sphere is green or red (Noack, 2016). Hopefully, they will thus be less likely to be hit by a BMW, which is unlikely to be virtual. As the doubling, tripling or even quadrupling of place becomes commonplace, the notion of privacy has become more layered, more complicated, more confused. As a result, it can be hard to know which norms ought to prevail. On the 373 bus, where I am in public in private in public, which privacy standards ought to apply? Presumably, the answer involves a complex layering of norms.³

My first point is that private and public are becoming increasingly *confused* by our internet interactions. To this I would like to add a second point, and that is the way that privacy is being *challenged* by our internet interactions. In recent years,

³ This complex layering of norms becomes even more complicated when we allow for the fact that the multiplicity of places in which I find myself are increasingly visible to others. This occurs, for instance, in the way my smartphone's location is being tracked, and in the way my internet use is being monitored. The seemingly private is often not actually private. This point is explored in detail below, in this chapter and the next.

this point has been repeatedly stated *in extremis*. “You have zero privacy anyway,” Sun Microsystems CEO Scott McNealy said way back in 1999. “Get over it!” (Rauhofer, 2008: 196) The idea is widely-held and oft-expressed: that privacy is an impossibility in the age of the internet. In 2010, Google’s Eric Schmidt responded to concerns that social media histories were going to hamstring people’s futures with the suggestion that they simply change their names and move on (Jenkins Jr, 2010). Along the same lines, legal scholar Jonathan Zittrain has proposed that we should have a mechanism for erasing our digital past and hitting the reset button: “As real identity grows in importance on the Net, the intermediaries demanding it ought to consider making available a form of reputation bankruptcy” (Zittrain, 2008: 228). Facebook founder Mark Zuckerberg agrees that people’s privacy is shrinking. For Zuckerberg, however, this isn’t a problem, because he says people’s norms are shifting too.

People have really gotten comfortable not only sharing more information and different kinds, but more openly and with more people. That social norm is just something that has evolved over time ... When I got started in my dorm room at Harvard, the question a lot of people asked was, ‘Why would I want to put any information on the internet at all? Why would I want to have a website?’ Then in the last 5 or 6 years, blogging has taken off in a huge way, and just all these different services that have people sharing all this information (Johnson, 2010).

These comments suggest that privacy is dead or doomed. As I argue throughout this thesis, privacy is still very much alive. It may even be perennial. However, it is under considerable pressure from our internet use. As early as 1968, a year before the internet appeared from the ether, Charles Fried noted that computer data storage had significant potential to compromise privacy (Fried, 1968: 485). Since then, the rise of the internet has made the challenge much greater (Reiman, 2004; Nissenbaum, 2010; Mayer-Schönberger and Cukier, 2013). As Mickle and Young write, “Convergent media make the invisible visible” (Mickle and Young, 2012: 129). On the internet, what was previously private is now regularly rendered public, and there is significant pressure on users to give up even more of what was once considered private. In simple terms, our internet use is challenging privacy.⁴

⁴ My phrasing here is deliberate. I am not claiming that the challenge is coming from the internet *per se*, but rather from the way we use the internet. I am thus departing from technological determinism. I will, however, go on to argue that there are values embedded in the internet and its platforms. Both these points are addressed in chapter two.

The challenge stems partly from the sheer quantity of data being shared by users. It also stems from the ease with which information flows on the internet. In 1997, James H. Moor invoked the phrase “greased data” to describe this ease. As Moor argued, “Given the ability of computers to manipulate information – to store endlessly, to sort efficiently, and to locate effortlessly – we are justifiably concerned that in a computerized society our privacy may be invaded and that information harmful to us will be revealed” (Moor, 1997: 27). Since 1997, the capacity of computers to store, sort and locate information has become more powerful by many degrees. In her research, danah boyd has identified four characteristics, or “affordances”, which shape today’s mediated environments and which impact heavily on privacy. These four affordances are: persistence, visibility, spreadability and searchability. Persistence, which ensures the durability of online expressions and content, means that those using the internet are “on the record” to an unprecedented degree. Visibility ensures there is potentially a huge, global audience of users who can bear witness. Spreadability explains the ease with which content can be shared, enabling people to mobilise for civil action with unprecedented speed, but also enabling malicious untruths to circulate as never before. And searchability means that online content, including esoteric interactions such as a glib Facebook post from 2008, tend to be easy to find (boyd, 2014: 11-14). In the past, people could rely on “privacy by obscurity”, but increasingly that option is being closed off to us (Zimmer and Hoffman, 2012: 175-176). The internet has changed the way we communicate, and changed the way our communications are shared, stored, filed and found. An oral conversation is easily forgotten; an email conversation is impossible to erase. Tellingly, emails are admissible evidence in legal proceedings, whereas hearsay is not. The former is regarded as reliable; the latter is regarded as problematic, due to the fallibility of human memory and the ethereal, contested nature of spoken exchanges. In this way and many others, as Bruce Schneier notes, it’s getting harder and harder for people to be ephemeral (Schneier, 2015: 128). This has been called the problem of “digital eternity” (Lindsay, 2014: 293-294).

The greased data described by Moor and the persistence, visibility, spreadability and searchability identified by boyd are two expressions of the same underlying principle: data is difficult to quarantine in the digital age. Research bears out this point. On social media, the “privacy leak factor” describes the way that users

reveal information not just about themselves, but also about their friends, often without realising. This leakage enables social networking services to create highly detailed profiles of its users, and even shadow profiles of non-users. This information can then be shared and sold (Sarigol et al., 2014: 95-96). If you have a smartphone, your carrier knows where you are at all times. It knows when you are at home, in a bar or at church. Given the same knowledge about other users, it also knows who is with you. In 2012, researchers analysed this data to predict where users would be 24 hours later, and were accurate to within 20 metres. As Schneier writes, “This is a very intimate form of surveillance” (Schneier, 2015: 1-2). Then there is facial recognition technology, which is far better at recognising people than people are, and has become highly adept at matching those people with their personal information. Already, the technology exists to build an app that recognises a stranger, then promptly calls up a summary of personal data (Acquisti et al., 2014: 13-15). Schmidt and Zittrain propose reputation bankruptcy; perhaps only face transplants will give users a chance of escaping their digital selves. Each of the examples I have just cited will be explored in more detail in chapter two, where I present a taxonomy of the challenges to privacy. For now, I merely wish to make the point that our internet use presents significant challenges to privacy. If we use social media, we are revealing not just ourselves, but our friends. If we use mobile phones, there is a good chance we are disclosing where we will be in 24 hours. And even if we simply exist in our own skin, our faces make possible the revelation of a great deal of personal information. On the net, then, we are exposing ourselves, and others, often without realising. In some cases, even the privacies of non-users are being exposed. In sum, our net privacy is shrinking.

In that context, let me turn for a moment to privacy itself. The concept of privacy, as defined and justified in chapters three and four, is slippery and elusive. This is especially evident amid the shifting interactions of the internet, among the unprecedented information flows made possible by mobile devices, social networking and data mining. However, the concept can be better understood if we make one significant preliminary distinction, between the *condition* of privacy and the *right* to privacy. The condition of privacy can be thought of as privacy *simpliciter*. It is the state of privacy, which is what I have, for instance, when I am home alone, not connected to the net, and unobserved by others. The right to

privacy, by contrast, concerns the situations in which I have some ethical or legal claim to privacy. The two are connected, but distinct. When I am home alone and not connected to the net, I presumably have a right to privacy too. Unless perhaps I am conspiring to commit a crime, whereupon legal authorities might justifiably be monitoring me. As Reiman writes, “I can have privacy without the right to privacy, say, when I successfully conceal my criminal activities. And I can have a right to privacy and not have privacy, say, when others successfully violate that right” (Reiman, 2004: 199). As such, let me qualify my claims. What I have been saying above is: first, that our internet use confuses privacy; and second, that it challenges our privacy. In these claims, what I have been referring to is the *condition* of privacy. I have been arguing, then, that our internet use confuses and challenges our *condition* of privacy.⁵ Indeed, to say that our condition of privacy has been confused and challenged in this age of webcams, cookies and government surveillance is hardly controversial. It is commonly accepted, I suggest, that people generally have less of the condition of privacy in the internet age than they had in the pre-internet age. This observation underpins statements to the effect that privacy is dead, and is supported by the research detailed above and below.

By contrast, I have not yet made any claims about the *right* to privacy. The fact that our internet use is confusing and challenging the *condition* of privacy does not necessarily impact our *right* to privacy at all. Admittedly, our internet use is, in all likelihood, having some effect. For a start, it is probably making the right to privacy more difficult to discern. Further, there does appear to be an emerging pressure on the right to privacy to adjust, by acknowledging and responding to shifts in the condition of privacy. My point, however, is that just because the condition has been diminished, that does not necessarily mean that the right has been diminished too. Nor does it mean that it should be diminished. Mark Zuckerberg says that people’s norms *are* shifting in favour of greater sharing and openness, implying that people are happily relinquishing some of their right to privacy. I will go on to argue that our internet use is putting pressure on our right to privacy but, contra Zuckerberg, that the right to privacy ought to remain intact

⁵ Throughout this thesis whenever the word “privacy” appears unqualified, I am referring to the condition of privacy; whenever I am referring to the right to privacy, I will explicitly spell out “right to privacy”.

and protected on the internet, partly in order that people recapture some of the condition of privacy they have relinquished.

This distinction between condition and right can be clarified with an example. Right now, I am not on the bus. Rather, it's a rainy Wednesday afternoon, and I am alone in my home office working at my computer. Here, I would appear to be in private. As I type these words, there is no one else in the room. People in other apartments might see me through the window, but they cannot read what I am writing. If I close the blinds, they will not be able to see me. *Prima facie*, the condition of privacy would appear to prevail. Further I seem to have a certain right to privacy. Social norms and legal prescriptions provide that strangers are not entitled to wander into my home. If I close the office door, my family tend to respect my privacy too. (It is worth noting that strangers are under a more onerous obligation not to interfere than my wife, children or Labrador. My right to privacy would seem to vary from stranger to kin.) In any case, my wife and children are out, so my privacy is almost complete. At least, my *physical* privacy is almost complete.

Meanwhile, however, I am connected to the internet as I work. Usually, apart from my word processing documents, I have several windows open at once, among them my email inbox, my Twitter homepage, miscellaneous news stories, various academic articles, and more. As I move between these windows, Gmail and Twitter ensure that I am, to some circumscribed degree, in public. I can, for instance, email a friend about a party before tweeting a quip about tonight's rugby league game. Here, then, the condition of privacy does not prevail, and my right to privacy is limited. For one thing, I know that Google scans my email inbox for data so that the company can then tailor my search results and target me with personalised advertising (Meikle and Young, 2012: 138). At the same time, I know that various websites I visit install cookies on my browser to facilitate my use of that website, and also to track my browsing habits (Schneier, 2015: 47-49). Further, under Australian law, my metadata is being stored by my internet service provider for two years, and various government agencies can access this metadata without a warrant (Scott, 2015). I also know that under the Five Eyes agreement, the governments of Australia, the US, England, New Zealand and Canada share information about their citizens (Schneier, 2015: 76). In these ways, my condition

of privacy and my right to privacy are limited in specific ways, even as I am protected by legislation including Australia's *Privacy Act*. Meanwhile, a hacker might be watching me through my webcam. This would be illegal and highly unlikely, but possible nonetheless, and would infringe both my condition of and my right to privacy, which is in part prescribed by the legislative prohibition on "computer intrusions" contained in the *Australian Criminal Code Act 1995*.⁶

Then there is my emailing and tweeting. When I email, perhaps if I tell my friend that my correspondence is personal and not for sharing, then I can have a strong expectation of an ethical, if not a legal, right to privacy. Perhaps this right might even be implied. I also have a certain legal right to privacy if my email contains a disclaimer that it is confidential and intended only for the recipient. With my email, I have a limited condition of privacy and a certain right to privacy; with regard to my tweet, however, I have neither. Twitter is a public forum, and hence any tweet I post can potentially be seen by millions of people. In short, when I am working on the internet in my home office, my virtual engagements occur in a complex web of private and public. My *virtual* privacy, therefore, is far from complete. Rather, I am simultaneously in private and in public, under a complicated set of norms and rules that apply regarding both the condition of privacy and my right to privacy. As I engage in the unremarkable task of working at my computer, I am, it turns out, in several places simultaneously, and various norms and rules apply all at once. My internet use means that private and public are knitted and knotted in a way that can be tough, if not impossible, to disentangle. In the internet age, privacy is confused and challenged.

Meanwhile, the corollary is that there is a simultaneous confusion and challenging of what is and what should be *public*. If it is late at night and you are in your bedroom, both the condition of privacy and right to privacy are conventionally presumed to prevail. If, however, you are also posting on Facebook via your smartphone, then you are simultaneously in the virtual forum of Facebook, which is, to some extent, public. Again, it is easy to imagine two theres colliding. Perhaps private norms will override the public, and your social media interactions will be stunted. More probably, public norms will prevail, and you will share

⁶ A government agent might also be watching me through my webcam. This might also be illegal, though a government agent is subject to different laws than a hacker. For further discussion of privacy laws including those relating to government surveillance, see chapters three, five and six.

liberally on Facebook. Indeed, research outlined in chapter two reveals that public norms tend to override private norms. Hence the public engagements of social media may well encroach upon and dominate the conventionally private space of the bedroom. Thus the public steps into the private. The reverse is sometimes true too, however, when the private steps into the public. This is evident in Moores's example of a woman having a phone conversation on a train, and so too in my example of being lost in the cocoon of a smartphone while riding the bus. Philosopher Wendy Brown noted the same phenomenon during a visit to Italy, observing a parade of Florentines in the Piazza di Republica talking into their handsets. The scene made her think of Hannah Arendt. As Brown writes:

In Arendt's view, the loss of clear demarcations between public and private imperils both. Yet the replacement of public conversation about shared matters of political, social, and economic life with individual cell phone conversations in the Piazza di Republica marks a diminution of public life next to which the presence of women with strollers in the public square (Arendt's lament) pales. Far from the cause of that diminution, such conversations, and above all our ready tolerance of them, are perhaps only its epitaph (Brown, 2004: 135).

For Brown, there has been a drastic diminution of public life, just as there has been an attendant diminution of private life. Together, Scannell, Moores, Arendt and Brown describe a world in which the boundaries between private and public have become confused and complicated, and the way in which, as a result, both spheres have become imperiled. Hence, as Vallor writes, "the integrity of the public sphere comes to look as fragile as that of the private" (Vallor, 2012: s. 3.1). Brown writes of an epitaph; but I propose that "private", like "public", still has meaning. Nonetheless, this blurring of boundaries is such a common feature of our internet use that often we don't even notice when it's happening. When we post photos from bed, or send an email from the Piazza di Republica, we rarely see the collision of places, let alone the casualties of these collisions.

II - The net: convergent, ubiquitous, multi-directional

To explore internet privacy, we first need to understand the various ways in which the internet is confusing and challenging privacy. In this section, I give a brief history of the internet before identifying three of its defining characteristics, which are combining to impact privacy dramatically.

First, what exactly is the “internet”, or “net”? According to Wikipedia:

The Internet is the global system of interconnected computer networks that use the Internet protocol suite (TCP/IP) to link billions of devices worldwide. It is a network of networks that consists of millions of private, public, academic, business, and government networks of local to global scope, linked by a broad array of electronic, wireless, and optical networking technologies. The Internet carries an extensive range of information resources and services, such as the inter-linked hypertext documents and applications of the World Wide Web (WWW), electronic mail, telephony, and peer-to-peer networks for file sharing (Wikipedia, 2016a).

Co-authored by the net’s users, this definition gives an insight into the nuts and bolts, the bits and bytes, the software and hardware that comprise the net.

However, what we want to uncover is what is *distinctive* about the internet, and in particular what is distinctive about the net in relation to *privacy*. To this end, in this section I examine three of the net’s defining characteristics: convergence; ubiquity; and multi-directionality. Together, these characteristics help to clarify how the net is confusing and challenging privacy.

In 1969, the internet creaked to life in the shape of ARPANET, built by the Advanced Research Projects Agency (ARPA) in the United States Defense Department (Castells, 1996: 6-7; Leiner et al., 2009; Meikle and Young, 2012: 29-32). With a series of breakthroughs, ARPANET soon drew in more networks and more users. In the early 1970s, the invention of email was enhanced with the addition of the CC function, thus supplementing one-to-one dialogues with multi-party conversations (Meikle and Young, 2012: 29). In 1973, the introduction of TCP/IP protocols enabled communication between different networks, not just within one network (Castells, 1996: 47-48). Then, from the late ‘70s, user discussion groups formed around disparate topics, including Usenet, and in 1985 the advent of the NSFNET, run by the National Science Foundation, enabled scientists to run programs on remote computers (Leiner et al., 2009; Meikle and Young, 2012: 29-30). The populist breakthrough, however, happened only once the net began to attract civilian users in large numbers, which came with the arrival of the world wide web, or “web”. In December 1990, Tim Berners-Lee activated the first website; then, in 1991, CERN (the European Organization for Nuclear Research) launched the web, complete with Uniform Resource Locators (URLs) to identify specific locations, HyperText Markup Language (HTML) and the HyperText Transfer Protocol (HTTP) (Berners-Lee, 1999). By marrying the

internet and hyperlinks, the web made the net - or at least a small part of it - easily navigable.⁷ Suddenly, new users started flooding in.

The dramatic effects of this new medium were immediately apparent, prompting claims of a digital “revolution” (Dutton, 1996; Castells, 1996). Others heralded the advent of the “network society”:

At the individual level the use of networks has come to dominate our lives ... Networks are becoming the nervous system of our society, and we can expect this infrastructure to have more influence on our entire social and personal lives than did the construction of roads for the transportation of goods and people in the past” (van Dijk, 2006: 1-2).

In this network society, wrote Manuel Castells, “both space and time are being transformed” (Castells, 1996: 376, 398). Regarding space, Castells argued that this would entail a rise in megacities and an increase in disparities between urban poles and their respective hinterlands (Castells, 1996: 380, 404); in terms of time, it would engender a breakdown in the rhythms associated with a lifecycle, which would involve people living as if age were immaterial and death does not exist (Castells, 1996: 446-451). More recently, however, theorists have argued that the promise of the networked era has not been fulfilled: along with market boom and hypersuccessful applications, there have occurred bust, viruses and spam (Zittrain, 2008: 238). Others have argued that it is inaccurate and unhelpful to speak of a digital revolution (Meikle and Young, 2012: 3).

Nonetheless, the internet *has* been transformative, including in the way it tends towards convergence, ubiquity and multi-directionality, three trends that are having profound impacts upon privacy. The internet’s tendency to “convergence” is revealed in its origins: the way it brings together networks that were previously discrete; the way it renders disparate hardware and software compatible; the way it connects people who were previously out of touch. Convergence is arguably *the* defining ingredient of the net specifically, and modern media generally (Meikle and Young, 2012: 2-5; Dwyer, 2015a: 14-17). Over time, the term’s meaning has grown. In the mid-‘70s, convergence denoted the coming together of the

⁷ The distinction remains significant: the internet is a network of networks that connects computers via hardware and software; the web, by contrast, is the system of *interlinked hypertext documents* that can be found on the internet via a web browser. The internet can be traced to the late ‘60s; the web was only launched in the early ‘90s as a way for users to access a portion of the internet. As I discuss below, the internet also houses the “deep web”, where users can be more anonymous and private than on the “surface web”.

computing, publishing and broadcast/film industries; in the '80s and '90s, it was used as a “techno-economic” buzzword, justifying a number of corporate takeovers and mergers in the IT, telecom, internet, media and consumer electronics industries (Lind, 2004). Since then, the term has expanded to encompass merging technology and content, and is now defined as, “the flow of content across multiple media platforms, the cooperation between multiple media industries, and the migratory behaviour of media audiences who will go almost anywhere in search of the kinds of entertainment experiences they want” (Jenkins, 2006: 2). At the heart of the current notion of convergence is the idea that content can move easily and efficiently across media (Nightingale, 2007: 28-29). On the net, content flows; data is greased, as Moor says. I can, for instance, watch my favourite TV show as it screens on free-to-air digital TV, or record it to my hard drive and watch it later, or watch it on the web via a TV station’s catchup service, or stream it via Netflix, or even buy a DVD box set.

In recent years, convergence has been joined by a related buzzword, “transmedia”, referring to narratives that break the bounds of one form and simultaneously exist in various forms: as movies; as comic books; as video games; as internet clips; as websites; as amusement park attractions; and so on (Jenkins, 2006: 93-130). Some of these expressions exist outside the internet (including comic books and amusement park attractions); but those that exist on the net are only possible thanks to the digital and networked nature of convergent media. Amid the phenomena of convergence and transmedia, content is breaking its bounds. In this way, privacy is greatly affected. Just as many stories can no longer be contained to one platform, one format, one medium, neither can personal data be contained to one platform, one format, one medium. The idea of convergence explains the way in which, on the internet, data tends to flow, and how that flow continues irrespective of whether data is private or public.

Above, I described the multiplication of place made possible by the net. Picture a woman who is physically in Berlin but who is simultaneously, thanks to her smartphone, in New York, the Himalayas and the virtual public(ish) forum of Facebook. All at once, the spaces that are New York, the Himalayas and Facebook (wherever that is) *converge* on this user’s mobile device. On the internet, the experience of place is thus being redrafted, and the effect on privacy

is dramatic. Often, privacy is tied to location. Different norms prevail in a café, an office, a toilet. When locations converge, it is no wonder that privacy is being so thoroughly confused and challenged, and that norms are difficult to discern. Meanwhile, there is a second significant effect of convergence that is more straightforward, and perhaps even more significant. That effect involves the way so much of our data is converging on the internet. Social media profiles are converging with smartphone locations are converging with text messages are converging with job applications, and so on. These bits and pieces are being linked with one another, and, what's more, are being linked with offline data (see chapter two). Convergence describes the way almost all data is wending its way onto the internet, which then has tremendous implications for privacy. Once all our data is in one place, it becomes much easier to find. The term "dataveillance" has emerged to denote the surveillance potential of so much accumulated, aggregated data (Vallor, 2016: 188-189).

The phenomenon of data convergence is all the more significant given a second defining characteristic of the internet: its ubiquity. Or, more accurately, its looming ubiquity. Simply, the net is playing an ever-increasing role in the lives of an ever-increasing number of people accessing it for an ever-increasing array of purposes. One measure is the amount of time users spend online, which is rising to such an extent that some psychologists now diagnose "internet addiction" via tests assessing pathological online use (Brand et al., 2014). A second measure is the total number of users, which has been growing at a dizzying rate since the web first made the net navigable in the early '90s: in 1995, fewer than 1 per cent of the world's population had an internet connection; by 2016, that figure stood at 47 per cent (UN, 2016: 6). This percentage is set to keep increasing. In 2015, Google and Facebook, who are traditionally rivals, revealed they are collaborating to connect the rest of the world's population using high-altitude balloons, solar-powered drones and other radical technologies (Simonite, 2015). A third measure is size. As at Monday, August 1, 2016, the indexed web contained at least 4.75 billion pages (2016). This is just the world wide web; the internet itself is much, much larger. Beyond the web is the "deep web", which houses content that is not indexed by search engines such as Google and is also known as "deepnet", the "invisible web" and the "hidden web". One study of the period 1984 to 2000 found that the deep web is 500 times the size of the surface web, and is growing

faster than the surface web (Bergman, 2001). The deep web also contains the “dark web” or “darknet”, which can only be reached with the use of an anonymous browser such as The Onion Router, or TOR (Rudesill et al., 2015: 7-10). Given that pockets of the deep web embrace anonymity, the exact size of the internet is unknown, and perhaps unknowable. Nonetheless, the net continues to grow dramatically as it plays an increasing role in the lives of a growing number of users. The net is, in short, tending towards ubiquity.

The trend towards ubiquity is accelerating with the advent of big data and the internet of things, which together mark out the current third age of the world wide web. In the 1990s, the web’s first incarnation was driven by commerce-based websites such as Amazon and ebay; then, in the 2000s, social media refashioned the internet to be participatory and collaborative, prompting the neologism “Web 2.0” (Meikle and Young, 2012: 65-68). More recently, as the internet has become more integrated in users’ lives, the terms “web3” and “Web 3.0” have been coined to describe the “ubiquitous computing web” (Gubbi et al., 2013). Tim Berners-Lee calls this the “Semantic Web”, which describes a web of data that can be processed by machines (Berners-Lee et al., 2001). At the heart of web3 and the semantic web are big data and IoT, or the internet of things. Big data involves the collection, storage and analysis of enormous, unprecedented quantities of information “to produce useful insights or goods and services of significant value” (Mayer-Schönberger and Cukier, 2013: 2). The potential is vast. In 2007 and 2008, mathematical modelling identified 45 Google search terms that coincided with the outbreak of the flu, giving the company the ability to discern, in real time, where and when outbreaks were occurring, and their severity. Previously, health authorities had only been able to identify outbreaks a week or two after they had started (Mayer-Schönberger and Cukier, 2013: 2). Meanwhile, the internet of things (or IoT) refers to the arrival of internet-connected cars, fridges, domestic appliances and more. It denotes the connection of physical things to the internet. Again, the potential is vast: internet-connected thermometers can monitor vaccines; moisture sensors in agricultural fields tell farmers of crops’ needs; and acoustic sensors in rainforests can help curb illegal logging (UN, 2015: 59). The IoT is growing dramatically: in 2000, about 200 million objects were connected via the internet; by 2020, an estimated 50 to 100 billion devices will be internet-connected (Perera et al., 2015: 32). Already, the United Nations has described the

combination of big data and the IoT as “the internet of everything and everyone” (UN, 2015: 60). Not surprisingly, as everything and everyone comes online, significant challenges to privacy have been identified (Perera et al., 2015: 34-38). As Dwyer notes: “The privacy implications of the ubiquitous Internet are quite literally changing how we live” (Dwyer, 2015a: 2). If our internet use both confuses and challenges privacy, then an internet that is ubiquitous (or that is approaching ubiquity) will only tend to heighten that confusion and those challenges.

One insistent symbol of the ubiquitous internet is the smartphone. Globally, more people now access the net from mobiles than desktops. In 2016, the UK media regulator found that “the smartphone is the preferred device for the majority of online activities” (Ofcom, 2016: 6-8). Mark Zuckerberg pinpoints 2013 as the year Facebook become a mobile-driven business (Dwyer, 2015a: 22). As handheld devices expand their capabilities, the phrase “mobile phone” has become obsolete, replaced by the more expansive “smartphone”, “mobile device” or “mobile media” (Goggin and Crawford, 2010: 224). Research is revealing the effect of mobile devices on crowdsourcing (Chatzimilioudis et al., 2012), health care (Putzer and Park, 2010) and tourism (Wang et al., 2012), *inter alia*. However, the biggest shift is more general: thanks to mobile devices, all users can now be *permanently* connected. Most of us no longer simply log on to the net when we reach our desk. Rather, we take the net with us wherever we go. As such, we don’t just consume media now, we *inhabit* the media world (Meikle and Young, 2012: 2). Or, as Gordon and de Souza e Silva write, “We don’t enter the web anymore; it is all around us” (Gordon and de Souza e Silva, 2011: 3). One estimate is that more than 90 per cent of people with mobile phones keep them within a metre of themselves 24 hours per day (Schmidt and Cohen, 2013: 172). Continuously connected, we can order groceries, send a flirtatious text, do the banking and pay our electricity bill with a few deft swipes and taps. As we do so, the distinction between user and media, between us and device, is beginning to disappear. In a study of the smartphone use of young Sudanese asylum seekers recently arrived in Australia, Evers and Goggin found that “mobile phones are not separate to bodies but part of them” (Evers and Goggin, 2012: 81). Tellingly, a term has emerged to describe the anxiety felt by those unable to use their mobile devices: nomophobia (Rauhofer, 2008: 185). Clearly, the rise of the mobile device

reveals the net's tendency to convergence (Dwyer, 2015b: 122-125). Even users and devices are merging. At the same time, the rise of the mobile device reveals how the net is being used more often by more people for more purposes. Portable, powerful and seemingly indispensable, mobile devices show the net's tendency to ubiquity, and are having tremendous impacts on user privacy (Dwyer, 2015b).

The blurring of the boundary between user and device is also evident in the emergence of wearable and embedded computing. Smartwatches are more compact and unobtrusive than smartphones and perform many of the same functions. Meanwhile, "activity trackers" are devices worn on the wrist to track steps walked, steps climbed, heart rate, sleep hours, calories consumed and burned, and more (Lanzing, 2016). Smartwatches and activity trackers allow users to *wear* the internet. In 2014, a poll found that one in six Americans own wearable technology, with numbers rising quickly (Nielsen, 2014). More dramatic are innovations that seek to *implant* the internet in users. In medicine, microchip implants are now used in heart pacemakers, in brain pacemakers to combat epilepsy, Parkinson's disease and depression, as well as in prosthetic knees and hips to provide data that aids rehabilitation (Michael and Michael, 2013: 78). Further, microchips are increasingly being implanted in human hands, wrists, forearms and triceps for non-medical purposes, including enhanced convenience and security for users and also the monitoring of criminals (Michael and Michael, 2013: 78-81). As bodies and bytes merge, passwords become passe, replaced by retinal scans, facial recognition, gait analysis, ear shape and voice recognition (Shankar et al., 2016). Meanwhile, virtual reality and augmented reality are on the rise.⁸ Inexorably, as wearables and implants become more common, as virtual and augmented reality become more widespread, users are becoming increasingly enmeshed in an ever-expanding net. The connection is constant; the net nears omnipresence; private and public blur.

A third characteristic marking out the internet is its multi-directionality, where each and every one of the net's users is a co-author, engaged in the creation of content, and thereby in the continual re-creation of the medium itself.

⁸ Virtual reality, or VR, enables users to be immersed in alternate realities via headset, hand controls and other devices. Increasingly popular among gamers, VR also has medical applications, such as in the treatment of stroke (Standen et al., 2011). Augmented reality involves superimposing digital elements onto reality, and can be used, *inter alia*, in medical contexts, by architects and engineers and for training soldiers (Behzadan et al., 2016).

Traditionally, media has been built on a model of few-to-many, in which a handful of media owners and editors curated content for the masses. This model prevailed in newspapers, on radio and on TV, where owners and editors decided which topics to cover, which angles to take and which viewpoints to endorse (Meikle and Young, 2012: 106). Significant decisions included what *not* to cover, such as racial minorities, niche causes or alternative opinions. The net changed this, however, making possible few-to-few communication, many-to-few communication and many-to-many communication. In 2003, journalism professor Jay Rosen described how the traditional media model had been upended: “The supremacy of the ‘one to many’ media system has ended, and vastly different patterns are emerging” (Rosen, 2003). Compared to what is now sometimes called traditional or mainstream media, the internet is participatory and democratic. It has the capacity to give voice to the voiceless.⁹ It is, in a word, interactive. All users can create content, and as they do so they are creating the internet. Blogs and social media allow users to post opinions, photos and observations. Traditional news outlets now rely on tip-offs, photos, videos and accounts provided by the general public. Digital technology gives users tools to create music, films and books, and the internet enables these works to be streamed, distributed and published. There has been a major shift from consuming audiences to creating audiences (Meikle and Young, 2012: 108). In 2006, Rosen captured this shift in the phrase, “the people formerly known as the audience” (Rosen, 2006). In the same spirit, Axel Bruns coined the term “produser” to describe the collaborative processes of content creation. “The very idea of content production may need to be challenged: the description of a new hybrid form of simultaneous production and usage, or produsage, may provide a more workable model” (Bruns, 2007: 99).

The collaborative potential of digital networks has also been described as “intercreativity” (Meikle and Young, 2012: 121-122). Indeed, Tim Berners-Lee says the internet was built by users solving problems and making things together (Berners-Lee, 1999: 182-183). In the net’s early years, each successive improvement was known as a “hack”, a term denoting a neat solution to a technological problem. In 1971, email was created after a hack of ARPANET’s

⁹ Well, to some of the voiceless. As I have noted, more than half the world remains disconnected. Moreover, the participatory, democratic ideals of the internet are not always realised, as I argue in the second half of chapter two.

messaging system (Meikle and Young, 2012: 29-30). A collaborative hacking spirit also underpins open-source software, or OSS, for which the copyright holder makes available the code and gives permission for others to study, change and share it (Meikle and Young, 2012: 121). And perhaps the best expression of the net's collaborative potential is Wikipedia, the "free encyclopedia that anyone can edit". As at April 2016, Wikipedia was one of the most-visited sites on the internet, with more than 5 million articles in English, millions more in German, Japanese and Russian, and more than 50,000 in Esperanto (Wikipedia, 2016b). On Wikipedia, user-generated knowledge flourishes in a participatory, philanthropic expression of multi-directionality.

Ray Kurzweil argues that rapidly accelerating advances in the science of computers, robotics and artificial intelligence will lead to a "technological singularity" in which mankind and machines will merge (Kurzweil, 2005). I won't go that far. Rather, I am merely arguing that the internet is becoming everything. Before the internet, chunks of data were discrete. Big data existed in one domain, as companies including Acxiom (see chapter two) compiled hard copy dossiers on individuals; CCTV existed in another domain, as businesses tried to discourage and catch shoplifters; and people's personal details existed elsewhere, such as in the filing cabinets of a doctor's office. Now, all these domains are being connected (see Vallor, 2016: 188-190). Thanks to the internet's tendency to convergence, all that data can be brought together. Thanks to the net's increasing ubiquity, more and more data is being collected. And thanks to the net's multi-directionality, that data can then be shared and sold. The point is not simply that more people are spending more time on their smartphones. The point is that people are putting more and more *of their lives* onto the net via their smartphones, and also via their laptops, their FitBits, their cars, and every other internet-connected device. Online is where a user's car registrations, insurance policies, health records, bill payments, banking transactions, jogging routes, student records, curriculum vitae, location details, grocery purchases, holiday bookings and more can be found. The net is where lovers sext, where children are automatically tagged in photos, and where traffic is assessed so that drivers can be told, as they slip behind the wheel, exactly how long their morning commute is likely to take. And often this data is being collected, sorted and stored by machines. The net is not just the web. Increasingly, the net is *all* our data. On

current trends, all data, or at least almost all data, is tending to find its way onto the internet. The effect on privacy is unprecedented. The combined effect of the internet's convergence, ubiquity and multi-directionality is that our world is becoming Panopticon 2.0.

III – Panopticon 2.0: A theoretical overview (and underview)

The internet's tendency to convergence, ubiquity and multi-directionality is having a deep impact on the condition of privacy. As you navigate the net on your smartphone, government agencies may be storing your metadata, companies might be tracking your clicks and a friend from your distant past might be uploading a photo to social media and tagging you as a pimply, awkward teenager. With privacy under threat, it is hardly surprising that scholars have compared life in the digital age to life in a panopticon (Reiman, 2004; Meikle and Young, 2012: 132-134). The metaphor fits like bespoke leg-irons.

In England in the late 18th century, Jeremy Bentham proposed a new model for a prison inspired by Panoptes, the mythical Greek guard with 100 eyes. He called it the panopticon, derived from the Greek words “pan”, meaning all, and “optikos”, meaning optic (TNSOED, 1993: 2081, 2085). After years of refinement, Bentham proposed a circular building with the guards' post at the centre and the prisoners' cells at the circumference. With each cell opening to the middle, the guards would be able to see through the bars and into every cell at all times. The design ensured that few guards were needed to watch many prisoners. The design also ensured that prisoners were potentially under surveillance at any moment, without ever knowing whether or not they were being watched. As Bentham wrote,

By blinds and other contrivances, the inspectors concealed (except in as far as they think fit to show themselves) from the observation of the prisoners: hence the sentiment of a sort of omnipresence. The whole circuit reviewable with little, or if necessary, without any change of place. One station in the inspection part affording the most perfect view of every cell, and every part of every cell ... (Bentham, 1811: 65)

On this model, prisoner privacy is severely limited; arguably, it is non-existent. And it seems that there are two losses happening here. In a Panopticon, the loss of free movement that attends all forms of incarceration is supplemented with the deprivation of privacy. The implication is that taking away privacy is a distinct

and compound punishment. Globally, the model was highly influential, inspiring prisons in Cuba, Holland the United States and elsewhere (Welch, 2013: 44).

Apart from influencing the world's prison builders, the panopticon model has also influenced philosophers, including Michel Foucault. In 1975, Foucault invoked the panopticon to explain not just prisons, but modern society generally. In capitalist societies, he argued, modern life is circumscribed by pervasive social control. The panopticon prevails literally in the design of jails, factories, schools, barracks and hospitals, but also metaphorically, in the way citizens have internalised mechanisms of control. Even if we are not under surveillance, Foucault wrote, we act as if we are. We have submitted to power. In a great irony, we have become the agents of our own giving-up-of-agency:

He who is subjected to a field of visibility, and who knows it, assumes responsibility for the constraints of power; he makes them play spontaneously upon himself; he inscribes in himself the power relation in which he simultaneously plays both roles; he becomes the principle of his own subjection (Foucault, 1977: 202-203).

Foucault's French title, *Surveiller et Punir*, literally translates as *Surveill and Punish*. The idea of surveillance was central to his thesis. And ever since Foucault, it has often been argued that modern technology has created a kind of *literal* panopticon, in which all our interactions are potentially under surveillance. The modern citizen is not free, it is argued, because, like Bentham's prisoner, she may be under surveillance at any time. She may still have freedom of movement, but not the freedom that attends privacy. This is the thinking that underpins David Brin's argument that humanity is inexorably headed towards a "transparent society". In 1998, Brin wrote that the only remaining issue is how adults will choose to live – to cooperate, compete and thrive – in such a transparent society. Will they opt for a surveillance state, in which police and governments watch over the minutiae of all citizens' lives? Or will they choose a society of participatory surveillance, in which everyone watches everyone (Brin, 1998: 3-9)? A more explicit link between digital technology and Foucault has been drawn by Jeffrey Reiman, who argued that a computer system designed to track vehicles might, as a side effect, build invasively detailed profiles of individuals.

If we direct our privacy-protection efforts at reinforcing our doors and curtains, we may miss the way in which modern means of information collection threaten our privacy by gathering up the pieces of our public lives and making them visible from a single point.

This is why the Panopticon is a more fitting metaphor for the new threat to privacy than, for example, that old staple, the fishbowl (Reiman, 2004: 196).

At the heart of the panopticon model is the notion that the many are made visible to the few, and hence that the powerless are kept under the control of the powerful. However, something is lacking when the metaphor is applied to the internet. Crucially, it doesn't take account of the net's multi-directionality, in which information potentially flows in all directions, rather than just one. One of the internet's defining characteristics, I have argued, is that it enables a participatory, collaborative intercreativity. In this regard, it is curious that Foucault's account omits any mention of the media. This is a perplexing omission, given that Foucault describes the emergence of panopticism in the eighteenth and nineteenth centuries, which was also a period of dramatic development in the media, marked by the arrival of newspapers, photography, telegraphy, telephony, recorded sound, cinema and broadcasting (Meikle and Young, 2012: 134). Sociologist Thomas Mathieson has addressed this omission, noting that while the panopticon allows the few to watch the many, the modern media enables the reverse, allowing the many to watch the few. On Mathieson's account, mass media, and particularly television, allows hundreds of millions to watch a limited number of celebrities, sportsmen, politicians and the like. Mathieson thus devised a reciprocal model: the *synopticon*, which refers to the many watching the few, and which draws on the Greek word "syn", meaning "together" or "at the same time". "Synopticism characterizes our society, and characterized the transition to modernity" (Mathiesen, 1997: 219). For Mathieson, the processes of panopticon and synopticon are operating simultaneously, working in tandem, with their reciprocal functions feeding on one another. "Together, the processes situate us in a viewer society in a two-way and double sense" (Mathiesen, 1997: 215).

Like Foucault, Mathiesen is seeking to expose the latent exercise of power in a democratic capitalist society, writing that, "Each from their side, like a pincer, panopticon and synopticon thus subdue ..." (Mathiesen, 1997: 230). My concern is elsewhere. Rather than power, I am concerned with the flow of information,

and with the question of who is watching whom, potentially and actually.¹⁰ In this sense, the internet bears out Mathiesen's theory. Certainly, the few are able to watch the many: hackers access strangers' computers; Facebook tracks users and non-users; government agencies such as the NSA seem the very incarnation of a modern Panoptes. On the net, panopticism prevails. However, synopticism prevails too, as the many watch the few. Millions of Kim Kardashian fans can follow the minutiae of her life via gossip sites, her Twitter and leaked sex tapes. The synopticon is also revealed in the notion of "sousveillance", a term which has emerged as the complement of surveillance. With its French etymology, surveillance suggests watching from above, or overseeing, while sousveillance suggests watching from below, or underseeing (Mann, 2004). As morsels of data are joined with more and more bits, we inhabit "a sousveillance society, one in which even the watchers are watched" (Vallor, 2016: 188).

The term has been used to describe WikiLeaks, the whistleblower website that uses collaborative software to enable information to be leaked anonymously. On a single day in January 2010, US army private Bradley Manning (now Chelsea Manning) downloaded 400,000 classified documents onto rewriteable discs while working as an intelligence analyst in Iraq. Manning then supplied them to WikiLeaks, who published them online. These became known as the "Iraq War Logs", and included "Collateral Murder", a redacted video shot in Baghdad in 2007 showing two Reuters journalists and 10 Iraqis killed by gunfire from a US Apache helicopter. Reuters had been trying unsuccessfully for three years to obtain the footage under *Freedom of Information* legislation; instead, WikiLeaks, styling itself as "the first intelligence agency of the people", posted the video to YouTube, where it amassed millions of views (Christians et al., 2012: 61). Similarly, in 2013 former National Security Agency contractor Edward Snowden leaked thousands of government documents to journalists, revealing, *inter alia*, the extent to which the NSA and other government agencies were engaging in global surveillance (see chapters two and six). In Snowden's case, then, sousveillance exposed surveillance. Synopticon revealed panopticon. And for both WikiLeaks and Snowden, sousveillance would have been impossible without the internet and digital data, which allowed hundreds of thousands of documents

¹⁰ That said, power is clearly an issue in a watcher/watched relationship. By knowing details of the watched's life, the watcher has power. In some cases, the watched, such as reality TV stars, have power too.

to be efficiently copied, shared and published. On the internet, users inhabit both panopticon and synopticon, with both expanding at an accelerating rate as they feed on one another (Mathiesen, 1997).

However, even a combination of panopticon/surveillance and synopticon/sousveillance is inadequate to describe the internet. What about the “co-veillance”, or mutual watching, of Fitbits and other self-tracking devices (Lanzing, 2016: 9; see chapter two)? How to explain the high quality audio, images and video that can now be captured on each of our smartphones, then summarily posted? And how about the user-generated site YouPorn.com, the quintessence of “intercreativity” and “produsage”, enabling people to upload pornographic home videos and to watch pornographic home videos uploaded by others (Meikle and Young, 2012: 139)? Multi-directionality is thriving; on the net, exhibitionism and voyeurism exist in volume and in symbiosis. The panopticon describes the NSA collecting data on US citizens; the synopticon describes US citizens keeping watch on the NSA; but FitBits, smartphone cameras and YouPorn.com suggest something more. They suggest the many watching the many and the few watching the few. We can think of this as the “banalisation of global surveillance” and the “democratisation of voyeurism on a planetary scale” (Virilio, 2002: 109). The panopticon evokes Orwell’s dystopian tyranny, with its slogan, Big Brother Is Watching You; the synopticon suggests a reversal, You Are Watching Big Brother. In the internet age, these reciprocal adages must be accompanied by others: Companies And Individuals Are Watching You; and its complement, You Are Watching Companies And Individuals. It is a world in which few watch many, many watch few, many watch many, and few watch few. The watching is top-down, bottom-up, and lateral. It is a world that has been described as omnipticon and as participatory panopticon (Mitrou et al., 2014: 1; Albrechtslund, 2008). As revealed by WikiLeaks and YouPorn, Ed Snowden and Kim Kardashian, data and privacy are being exposed in an unprecedented manner, and in all directions. In the internet age, we are all potential Panoptes.¹¹

To this proliferation of watching, there is one more significant point to add, concerning the proliferation of recording. The sketch I have drawn so far,

¹¹ The word “potential” is key here. As I detail in chapter two, I do not mean to suggest that information *actually* flows without limits. Imbalances remain; limits exist.

allowing for surveillance, sousveillance and lateral watching, addresses boyd's affordances of visibility and spreadability (boyd, 2014: 11-14). It does not, however, sufficiently address boyd's two further affordances of persistence and searchability. The internet means that data is often permanently recorded and conveniently accessible. This means that each individual, company and government is not just a potential Panoptes. After all, Panoptes may have had 100 eyes with which to watch his neighbour, but he was unable to sift through all his neighbour's correspondence with the help of a keyword search. He was unable to view footage of what his neighbour had done yesterday, or last week, or last year. And he was unable to aggregate thousands of bits of information captured over many years into a coherent, accurate profile, including data supplied by other people, and data inferred from other data. On the internet, permanence prevails. Users leave traces. And those traces are often accessible. The internet, structured around links and associations, adeptly leads users to relevant data. So yes, Big Brother Is Watching You, but remember also that Big Brother Was Watching You - And Recording It. If your current online activity can be accessed, so too can your past online activity. This makes for remarkably efficient surveillance. As Schneier writes:

On the internet, surveillance is ubiquitous. All of us are being watched, all the time, and that data is being stored forever. This is what an information-age surveillance state looks like, and it's efficient beyond Bentham's wildest dreams (Schneier, 2015: 32).

Thanks to the internet, individuals, companies and governments can all not just watch one another, but record one another, then search efficiently through those recordings. This means that we are all potentially Panoptes, each with 100 cameras on our 100 eyes. We are watching, and recording too. In web3, we can see more of the present, and also more of the past, which then gives us a certain ability to predict future behaviour. Welcome to life in Panopticon 2.0.¹²

¹² Again, the word "potentially" is key. In chapter two, I argue that there are in fact limits on watching; and in chapter six I argue that these limits on watching ought to be clearly articulated and protected by a series of legal and extra-legal measures. Panopticon 2.0 is a theoretical model that seeks to bring into relief current trends. It is not intended to prompt despair and apathy, but rather to motivate privacy advocates and policymakers into action.

Conclusion

It isn't just on buses that people can be found absorbed in screens, lost in their own little worlds. It happens in public places, from parks to restaurants, from beaches to shopping malls; and it happens in private places too, such as bedrooms and bathrooms. And when we fixate on our screens, we experience the multiplication of place. Suddenly, we are in two, three or more locations at the same time. Some are physical; some are virtual. Some are private; some are public. And sometimes these places butt up against one another, leading to a confusion of spaces, and thereby to a confusion of public and private. The *condition* of privacy is being confused and challenged by our internet use, making the right to privacy hard to discern. In the process, more and more user data is finding its way onto the net, with dramatic implications. Location data can reveal where users are likely to be in 24 hours; software can recognise faces and match them with personal information; social media can construct "shadow profiles" of non-users. This confusion and challenge to what is private (and simultaneously to what is public) I have attributed to three of the internet's defining characteristics: convergence; ubiquity; and multi-directionality. First, the internet is converging. Technology, users and content are coming together. Second, the internet is approaching ubiquity, particularly with the emergence of big data and the internet of things, as well as mobile devices and wearable computing. It is becoming the internet of everything and everyone. And third, the internet is multi-directional, built on a collaborative, participatory back-and-forth in which content is co-authored, and in which the very net is being continuously rebuilt by its users. The internet is a work-in-progress, the apex of intercreativity. In turn, these attributes led me to advance a theory of Panopticon 2.0, in which all potentially watch and are watched.

The internet depicted in the film *Her* is one vision of how a more advanced version of the internet might affect privacy. At one point, Theodore is surprised by how quickly Samantha can learn all about him. "You mind if I look through your hard drive?" she asks, before taking just seconds to sort his emails and his contacts. Soon she knows him inside out. At the same time, Samantha is also the operating system for thousands of other users. She knows them inside out too. *Her*

provides one account of the way a future internet might confuse and challenge privacy. Samantha is a kind of Panoptes. She knows everything, or almost everything, about Theodore and thousands of other users. Theodore, however, is not like Panoptes. He knows nothing about the other users (aside from the fact that they exist), just as they presumably know nothing about him. This world then is not Panopticon 2.0. Furthermore, *Her* does not align with my account of the internet as tending to multi-directionality. Certainly, there is convergence (between user and net) and ubiquity (Samantha is *everywhere*), but their relationship is one-sided rather than bi-directional. Samantha is party to Theodore's privacies, but he is not privy to hers. As a result, there is a gap between the film *Her* and the model of Panopticon 2.0 I have been sketching.

Similarly, there is a gap between Panopticon 2.0 and today's internet. This gap is what I now wish to explore. My argument has been that we are tending towards Panopticon 2.0. However, Panopticon 2.0 remains merely theory. Omniscience remains the net's potential; in reality, there exist limits and checks on the flow of data and the sharing of privacies. Alongside convergence, there are also signs of divergence, including in the proliferation of digital devices. Ubiquity is not yet complete. There are blocks on multi-directionality. What's more, these limits and checks discriminate. There are imbalances in the flow of data. In theory, individuals, Facebook and the NSA can potentially access all the internet's data. In reality, Facebook and the NSA are closer to actual Panoptes than the average individual. In the next chapter, I construct a taxonomy of the challenges to privacy to show in more detail how today's internet diverges from Panopticon 2.0, even as we tend towards it. I then ask what, if any, ethical norms prevail on the net. The trend, I have been suggesting, is towards Panopticon 2.0. In the next chapter I ask: are we nearly there yet?

Chapter 2

A triple threat and an epic clash

Minority Report is a sci-fi thriller starring Tom Cruise as John Anderton, a cop who arrests people before they offend. In the year 2054, Anderton and his “Precrime” department have been so successful that the murder rate has dropped to zero. It’s all thanks to the predictions of mutated humans called “precogs”. In this respect, the film is pure fantasy. Despite significant advances, police departments are not, according to my research, on the verge of employing infallible mutant soothsayers. Even so, the film was eerily prescient. In some jurisdictions, as I discuss below, a more generalised form of predictive policing has already been implemented, using big data in much the same way John Anderton used precogs. Equally prescient was the film’s depiction of personalised advertising. As Anderton moves about public spaces, he is met with a babble of voices hoping to attract his attention. These voices come from huge personalised adverts for cars, beer and credit cards. The ads are like TV spots, only tailored to each passer-by. “John Anderton,” says a man insistently from a screen depicting an enormous beer. “You could use a Guinness right now.” On another screen, a woman is spruiking a car. “Lexus,” she says. “The road you’re on, John Anderton, is the road less travelled.”

In the previous chapter, I showed that the internet is confusing and challenging the condition of privacy, in part due to three of the net’s defining characteristics: convergence; ubiquity; and multi-directionality. I argued that a fitting model to describe the net is not the Panopticon of Bentham and Foucault, but Panopticon 2.0, in which everyone can see everyone in the present, in the past and even, potentially, into the future. In Panopticon 2.0, surveillance is joined by sousveillance, and by further lateral watching. Finally, however, I suggested that in fact there *are* limits on this watching. In theory, everyone can watch everyone; in practice, there are restrictions. Each of us is merely a potential Panoptes, not an actual Panoptes. In this chapter, I interrogate the gap between potential and actual to see just how all-seeing the internet actually enables us to be, before then investigating some of the ethical values and viewpoints that prevail on the internet to explore some of the normative influences affecting online behaviour.

The chapter is divided into two sections. The first section seeks to classify challenges to individual privacy by sorting them into three categories of threat: from individuals; from companies; and from governments. The threats from these three distinct sources, I argue, necessitate distinct responses. In this section I ask, *inter alia*, just how lifelike and probable are the personalised advertising and predictive policing of *Minority Report*. I also describe various forms of pushback against these threats to privacy, thereby seeking to detail the limits to surveillance, sousveillance and other forms of watching. The chapter's second section turns to an examination of some of the ethical principles and approaches that prevail on the internet. Here I argue that some of the net's champions argue that conventional ethics do not apply on the net, and ought not apply. I respond that norms, conventional and otherwise, already are being applied online. Some of these norms are user-imposed; others are embedded in the very architecture of the net and its platforms. The internet, it turns out, is neither norm-free nor norm-neutral. Rather, it is an ethically-contested space. Limits to watching do exist online, including normative limits. Against this background, I then propose that norms *should* prevail on the net, including norms pertaining to privacy. Despite the arguments of anarchists, and contra the arguments of technological determinists, I argue that competing rights and freedoms need to be considered before it is possible to decide in which cases privacy ought to be protected. As it happens, this is precisely the point that viewers are left to contemplate at the conclusion of *Minority Report*.

I – A privacy trinity: three layers of encroachments

Acxiom has been described as “one of the biggest companies you’ve never heard of” (Rushkoff and Goodman, 2004). There is, however, a good chance it knows you. Indeed, it probably knows you well, right down to your age, race, sex, weight, height, marital status, education level, politics, shopping habits, health issues, holiday plans, and much more (Behar, 2004; Singer, 2012; Bambauer, 2013). In the booming information economy, personal data has been described as the “new oil” (see Dwyer, 2015b: 129). Companies such as Acxiom are “new oil”

miners. They are, in the words of the US Federal Trade Commission, the “unseen cyberazzi who collect information on all of us” (Singer, 2012).

As we have seen in chapter one, users are sharing more and more information online, both about themselves and others, both wittingly and unwittingly. Companies such as Acxiom are aggregating all that information. A data mining success story, Acxiom has dossiers on more than half a billion internet users worldwide, with an average of 1500 “data points” on each consumer (Singer, 2012). The process is straightforward. Each user is given an identifying number, and to this number Acxiom appends information according to location, credit card transactions, hobbies, interests and other markers. Unobtrusively, it accumulates this information in the background, then sells it. This makes for a very profitable business model: in 2012, Acxiom’s revenue was more than \$1billion (Bambauer, 2013: 667-668). Many similar companies abound, including Experian and Equifax (Mayer-Schönberger and Cukier, 2013: 100). Often the data they hold is sold; sometimes it is stolen. Accordingly, data held by Acxiom has been accessed by hackers (Bambauer, 2013: 668), social media companies (Sengupta, 2013) and government agencies (Antón et al., 2004). In other words, if Acxiom knows a lot about you, a long list of other people, organisations and agencies know a lot about you too, most probably without you ever noticing.

Challenges to privacy come in many forms. Some are deliberate, some are incidental. Some are trivial, some are concerning. Some are justified, some are illegal. To begin to understand in more specific terms the nature of these challenges, it helps to categorise the various challenges to individual privacy into categories based upon the *origin* of that challenge.¹³ The example given above suggests that a large number of very different challenges can be traced back to a single source: Acxiom. By dealing in user data, this company challenges user privacy. As such, Acxiom would logically be subject to the same ethical and legal responsibilities as all other companies. However, Acxiom is not acting alone. In a

¹³ Note that I am talking about individual privacy, which is the focus of my thesis. Other types of privacy, including the privacy of a couple, of friends, of a family, are also possible, but are not my concern. However, my account of individual privacy will build towards a notion of *relational* privacy, in which individuals are recognised as beings-in-relation. It is also worth noting that my tripartite model is not unprecedented. *Inter alia*, Schneier (2015) draws a clear distinction between government and corporate challenges to privacy; while Andrews et al. (2015) separate individual, corporate and government threats in their analysis of remote access of webcams.

broader sense, the challenge to privacy outlined above comes from a number of sources. Various actors can harness information aggregated by Acxiom: hackers can and do illegally access users' personal data; other companies can and do use it to build up accurate user profiles; and government agencies can and do supplement these profiles with information of their own. Hence I am suggesting that challenges to privacy can come from three sources: from individuals (including hackers); from companies and organisations (including Facebook and Acxiom itself); and from governments and public institutions (including the National Security Agency and the Australian Tax Office). Further, I am suggesting that, thanks in part to companies such as Acxiom, sometimes the source of the threat can be confused and blurred.

Note, however, that this tripartite taxonomy is merely a rough guide. There are tough in-between cases. How do we classify a non-government organisation (NGO) such as Greenpeace or the Electronic Frontiers Foundation? What about an organisation such as WikiLeaks? Even trickier cases arise whenever private enterprise and government are mixed together, as with some prisons and police departments. In chapters five and six, I argue that there are certain situations in which governments are justified to encroach upon privacy but companies are not. Does this mean that privately-run police departments, say, should be required to abide by the information-handling standard set for governments, or rather by the stricter standard set for companies? These questions may be complicated by a point raised above: that challenges to privacy can sometimes come simultaneously from more than one source. I leave it for others to flesh out this system of classification more fully, or otherwise to suggest a better alternative. I propose simply that this three-pronged system covers most cases clearly and well. While there is a significant overlap between these three challenges, there is nonetheless something distinctive about each, both in the sources' capacity to gather data (as I discuss in this chapter) and also (as I discuss in chapters five and six) in the rights and responsibilities that ought to apply.

i. The threat from individuals

In my proposed three-pronged structure, a first challenge to privacy comes from individuals. This can occur in several ways. In one category of cases, there are

deliberate and willful intrusions, often by hackers. Such invasions may be criminal, and can lead to jail terms for those responsible (Gander, 2014; Yuhas, 2016). In a second category of cases, surveillance is conducted not by hackers, but by parents and partners. In a third category, including revenge porn and sext-forwarding, it is *on-sharing* that constitutes a potential invasion.

The case of Cassidy Wolf falls into the first category. In 2013, shortly before she was named Miss Teen USA, the Californian teenager received an email from a stranger. Opening the message, Wolf was surprised to discover it contained naked photos of herself, taken surreptitiously in her bedroom via her laptop's webcam. Her laptop, it emerged, had been hacked using malware. After a police investigation, the emailer was unveiled as a classmate who had been spying on Wolf for a year. "It was traumatising," Wolf said later. "It's your bedroom. That's your most private, intimate space and that's where you should feel the most safe" (Perez et al., 2014). The teenage hacker who spied on Wolf had been spying on a total of 100 to 150 women. There are many like him. According to one estimate, between 2013 and 2015 more than 100 million people were notified that they had been victims of a data breach, thus exposing data held by retailers, health insurance companies and entertainment companies, among others (Chideya, 2015). Social networking sites are particularly vulnerable to malware risks (Mansour, 2016). Other privacy-compromising hacks involve "phishing", an attempt to get sensitive personal information such as passwords by sending emails purporting to be from legitimate sources (Guthrie, 2016). Another emerging internet practice is "doxing", or "doxxing", a type of public shaming which involves publication of someone's name, address, work details and other sensitive details, followed by online and offline harassment (Trottier, 2016). Doxing victims can be people who offend online communities, such as Brianna Wu, a game developer who challenged misogynist bullies online and then, after her address was published, received death threats (Stuart, 2014).

A second category of intrusions by individuals involves not strangers, but persons familiar. Surveillance apps have been developed to enable parents to keep track of their children by monitoring every text sent, number dialled and website visited, as well as revealing a child's location at all times and everything she posts to social media. The TeenSafe app boasts that parents will be able to "find the way

to their child's mind", and the company's chief executive argues that the app is legal and legitimate: "It's absolutely legal for a parent to do this discreetly ... What we believe is that when it comes to protecting your child from these things - privacy is trumped by protection" (Morris, 2015). Surveillance apps can also be used by adults to keep track of other adults. Employers can track employees; wives can track husbands; hospitals can track patients. Already, spyware has featured in criminal trials, including the conviction of Simon Gittany for the murder of his girlfriend Lisa Harnum in Sydney in 2011. By using an app named MobiStealth to track Harnum's messages, Gittany learned she was planning to end the relationship (Olding, 2014). Researchers are finding such technologies are increasingly being used to stalk and harass women in the context of domestic violence (Woodlock, 2016: 24-28).

A third category of cases involves the *on-sharing* of private materials. "Sexting", the sending and receiving of sexually explicit material, is increasingly common, with one survey finding that 88 per cent of people aged 18-82 say they've sexted in the past year (Stasko and Geller, 2015). It is also common among teenagers (Lee et al., 2015: 5; Strassberg et al., 2013). Privacy issues can arise when a person willingly shares material, only for it to be further on-shared without that person's knowledge and/or consent. This is known as "image-based abuse" and it is rife (Henry et al., 2017). One Australian study found 20 per cent of teens had shown a sext to someone else, and six per cent of teens had forwarded a sext to a third party for whom the image wasn't intended (Lee et al., 2015: 42-43). The consequences can be serious, including expulsion from school for the victim (Lenhart, 2009: 15-16). Similarly, "revenge porn", addressed in chapters five and six, involves the non-consensual sharing of images that were initially shared consensually.

The extent of the challenge from individuals is also linked to the emerging trend of sousveillance, as described in chapter one. In April 2016, the "Panama Papers" were described as history's biggest data leak, involving the unauthorised release of 11.5 million documents from a Panamanian law firm detailing the offshore tax arrangements of more than 200,000 entities, including political leaders (Harding, 2016). In 2013, Google Glass arrived with much hype and expectation. The wearable technology was unobtrusive, resembling regular glasses with a small

clip attached to one side, enabling users to go about their daily business while simultaneously engaged in online activities, including surreptitiously livestreaming from an in-built camera (Meese, 2015). As one technology writer noted, “To look at a person wearing Glass is to look at a camera staring back at you. It puts you in a state of uncertain surveillance, never sure whether or not you are being recorded” (Price, 2014). Even though Google Glass flopped (seemingly over privacy concerns), the development of sousveillance-facilitating technology is ongoing, conducted largely in secret (Metz, 2015: see chapter four).¹⁴

ii. The threat from companies

In 2013, whenever I logged into Gmail, my inbox displayed a list of recently-received messages. Above those emails sat a small, unobtrusive one-line advertisement. This ad consisted of about a dozen words, complete with a hyperlink to an external website. Gmail is the email system run by Google, and the content of this advertisement, as Google’s own information explained, was determined by the content of my inbox, as well as by my other online activity. In this way, I then received advertisements for higher education opportunities (no doubt drawn from my correspondence with various universities), for theme parks (presumably from my emails about my kids) and vanity units (probably from an email I once sent a plumber about bathroom fittings). In 2013, this was the apex of ads: subtle; responsive; personalised. By 2017, this model had long since been abandoned and superseded, as Google constantly updates and refines its native advertising (Lardinois, 2015).

In this section, I explore the challenges from companies, including: the ability of social media services to create shadow profiles of non-users; the way offline data is being merged with online data; the phenomena of data mining, big data, machine learning and the internet of things; the way social media reveals users’ friends; the challenges of wearables and cookies; developments in personalised advertising; and the shift to cloud computing.

¹⁴ I am not claiming here that sousveillance is wrong simply because it can invade privacy. The sousveillance involved in exposing race-based police brutality in the US under the hashtag #blacklivesmatter, often with footage shot on smartphones, is a case where “underseeing” has revealed wrongdoing and precipitated justice. *Prima facie*, sousveillance challenges privacy; but, as I argue in chapters four, five and six, the right to privacy must always be balanced against other rights and freedoms.

The challenge from companies is, perhaps, the most pressing challenge of all.

The overwhelming bulk of surveillance is corporate. We accept it either because we get value from the service or because we are offered a package deal that includes surveillance and don't have any real choice in the matter (Schneier, 2015: 47).¹⁵

One result of such surveillance is personalised advertising. A more extreme result comprises the shadow profiles maintained by social media companies, which challenges the privacy even of those who do not use social media. If you have a social media account run by company x, then x probably knows a lot about you. However, even if you *don't* have a social media account run by company x, then x probably still knows a lot about you. What's more, all this knowledge is not just likely to include data you have chosen to share, but data you have *not* chosen to share. That is because data has been obtained from other sources (such as your friends), has been deduced from the data you have chosen to share, or has been inferred from the person you are and the company you keep. Online, this phenomenon has been captured in the phrase "privacy leak factor" (Sarigol et al., 2014). In these ways, social networking services are able to create full "shadow profiles" for people who do not have an account (Sarigol et al., 2014; Elmer, 2015). Shadow profiles have received scant academic attention; but in 2013, after a security vulnerability was exposed, evidence emerged that Facebook does create shadow profiles on as many users as it can (Blue, 2013; Sarigol et al., 2014: 95-96). As researchers concluded: "not having an account in an OSN [online social network] does not guarantee a higher level of privacy, as long as one has enough friends who already are in the OSN" (Sarigol et al., 2014: 104). In 2015, three men from Illinois filed lawsuits seeking damages for these shadow profiles (Zara, 2015). In May 2016, a California judge ruled that the case could proceed (AAP, 2016a). Court cases concerning shadow profiles are taking place in other jurisdictions too. In June 2016, a Belgian court indirectly gave Facebook authority to track non-Facebook users, as well as Facebook users who aren't logged into the service (Anthony, 2016).

Companies that operate on the net tend to want to know about their users, and about their non-users. The more companies know about their users, and about

¹⁵ My account of the threat from companies and governments is largely informed by Schneier (2015).

non-users, the more likely they are to profit and thrive (Elmer, 2015; Schneier, 2015: 46-47). In the information age, data is currency.¹⁶ Above, I outlined how a data mining company such as Acxiom builds detailed dossiers on individuals. After the attacks of September 11, 2001, the US government, frustrated in its efforts to find information about the terrorists, released the names of 19 hijackers in a bid for help. Acxiom located 11 of them in its databases, providing the FBI with current addresses, former addresses and names of associates, among other details (Behar, 2004). Alongside data mining companies, companies such as Google and Facebook collect their own data on users. Various companies then enter into deals to share data, as Facebook has with Acxiom, Experian and Quantum (Baker, 2015). This shift to user-as-product has been a feature of the social media-led internet which has emerged since 2000 (Meikle and Young, 2012: 65-68). In the “Web 1.0” of the 1990s, companies sold products to users; in Web 2.0, the business model involves aggregating great quantities of data about users. As Schneier writes, “The primary goal of all this corporate Internet surveillance is advertising” (Schneier, 2015: 47). With all this data, companies are better able to advertise to increasingly specific demographics, if not individuals. Forget blunderbuss marketing; the new model relies on identifying adventurous 20-somethings with a large disposable income who like to eat Japanese. The business model is partly predicated on a loss of privacy. As the internet has matured, there has been a deep shift. Once, users were sold products; now users are sold *as* products (Meikle and Young, 2012: 67).

Social media, or social networking services, are particularly well-placed to gather user data, enabling a form of “participatory surveillance” (Albrechtslund, 2008; Strassberg et al., 2013). Areas of concern are manifold, including: the availability of user data to third parties; the ability of facial recognition software to identify people in photos; settings that are public by default, private by effort; the use of cookies to track users, even when users are off-site; the use of location data for illicit monitoring; the sharing of data with government agencies; and the pressure on users to share material imprudently (Vallor, 2012: s. 3.1). The biggest social media platform is Facebook, which launched in 2004, and which by 2016 counted 1.65 of the world’s 7.4 billion people among its users (Statistica, 2016). This

¹⁶ Unless, say, the company’s very existence is predicated on respecting privacy and *not* sharing such data, as is the case, for instance, with initiatives such as the privacy-protecting search engine DuckDuckGo.com, discussed below.

despite the fact that it is banned in the world's most populous country, China. What's more, studies have revealed that social media users feel pressure to share. Many users accept as "friends" other users they do not know, and who will thus have access to their names, birthdays, photos and many other personal details (Debatin et al., 2009).

In chapter one, I argued that the internet is tending towards convergence. One expression of this trend is "data convergence", which describes the way social media platforms are actively seeking to bring together as much data as possible about all their users (and, it seems, non-users). This involves deleting fake accounts. As a spokesperson said in 2012:

It's something we monitor vigilantly. We want to ensure that one of the core tenets of Facebook is that you have your unique identity on Facebook ... We have an advantage because we are a true identity platform so we can quickly figure out if anyone is their true self on Facebook ... What we are looking for is people who have widespread fake user ID accounts to make sure we take them out of the system. We call them bad actors ... (Edwards, 2012).

Similarly, in 2016 Twitter launched a campaign to "verify" its users (McGoogan, 2016). Data convergence also involves supplementing online data with offline data. In June 2016, Facebook contracted with data analytics firm Quantum to measure the impact of advertising on in-store sales, just as Pinterest signed with Oracle Data Cloud to gauge the offline shopping habits of those who had seen their "Promoted Pins" (Canning, 2016; Slefo, 2016). Offline and online will continue to merge with dramatic advances in facial recognition technology, which has seen an upsurge in the number of publicly available facial digital images accompanied by the fast-improving ability of computers to recognise the faces in them (Acquisti et al., 2014: 1). Facial recognition technology is already better at recognising people than people are, and, once a person is identified, personal information is easy to attach (Acquisti et al., 2014: 4, 10-12). Thanks to the documented pressure to share, the merging of offline and online, advances in facial (and other) recognition and attempts to weed out "bad actors", an increasing amount of personal data is being attached to real, identifiable people.

Data mining and machine learning pose particular challenges. Already, data mining is used by a wide range of industries, including human resources. The San Francisco company Gild scours online data to find potential job candidates for

specific companies; another Californian company, Visier, uses mountains of data to tell companies which of its employees are at imminent risk of quitting; and Wall Street firms including Goldman Sachs have invested in Digital Reasoning System, which analyses billions of emails, phone calls and online chats to predict which employees are likely to behave illegally (Sklar, 2015). Meanwhile, machine learning, the development of computers that can teach themselves, is making the collection and sorting of information wildly more efficient. Google's artificial intelligence program, "DeepMind", is working to mimic the visual cortex of animal brains to perform complicated analytic tasks (Turner, 2016). Increasingly, the internet is becoming all our data, and that data is being used by employers, insurance companies and social networking services. Together, data mining and machine learning, combined with the advent of the internet of things, are creating a perfect storm of data convergence, with privacy directly in its path.

Facebook and the internet tend to encourage users to share information about themselves *consensually* (Debatin et al., 2009). This seems unproblematic.¹⁷ More problematic is that Facebook and the internet tend to facilitate users sharing significant amounts of information *nonconsensually* and *unwittingly*, both about others and themselves. This point is brought into relief by the very existence of shadow profiles, as discussed above. Another example is in the way I can upload a personal photo of another user to Facebook and "tag" them as I do so. By "tagging", I am identifying that person and creating a link that other users can follow to learn more about them (Facebook, 2016b). Researchers have shown that tagging on Facebook can reveal highly sensitive user attributes (Pesce et al., 2012). On Facebook, I can unwittingly reveal personal details about others. At the same time, I can unwittingly reveal personal details about myself. In a 2013 study of 58,000 Americans, psychologists found that a Facebook user's "likes" reveal an uncannily accurate personality profile, including sexual orientation, racial heritage, political leanings, drug use and intelligence level. For instance, "likes" that predict low intelligence include Harley Davidson motorbikes and the band Lady Antebellum, while "likes" that predict male heterosexuality include

¹⁷ I say "seems" because the issue of consent is, alas, not so straightforward. Ostensible consent is not always morally justifying consent. What's more, as we have seen with sexting and revenge porn, data that is shared consensually can then be re-used in ways for which consent has not been given. The issue of consent is the subject of chapters five and six.

basketballer Shaquille O’Neal and “being confused after waking up from naps”.

The researchers concluded:

Commercial companies, governmental institutions, or even one's Facebook friends could use software to infer attributes such as intelligence, sexual orientation, or political views that an individual may not have intended to share (Kosinski et al., 2013: 5804-5805).

A later study confirmed these results (Mansour, 2016). Then there is “homophily”, which describes the tendency of people to befriend others with similar traits (Mislove et al., 2010: 259). In a 2010 study of two online social networks, researchers found users tend to be friends with others who share their attributes, and that communities form around users that share attributes. By applying an algorithm, the researchers then found that, with as few as 20 per cent of users providing attributes, they could often infer the attributes of the remaining users with great accuracy. As the researchers note, users who wish to remain private need not just keep their attributes private, but need to ensure that their list of friends remains private so that those attributes cannot be inferred (Mislove et al., 2010: 260). On social media as in real life, birds of a feather flock together, with tremendous implications for companies offering, say, health insurance or credit cards (Mayer-Schönberger and Cukier, 2013: 92).

Meanwhile, companies including Fitbit and Strava make wearables that enable users to track steps, sleep, calories burned and more. In 2016, this growing industry was worth an estimated \$700million annually, with its users known as “self-trackers” (Lanzing, 2016: 1, 9). The most dedicated self-trackers are members of the “quantified self movement”, whose motto, “self-knowledge through numbers”, reads like a modern extreme of the Socrates dictum, “the unexamined life is not worth living” (quantifiedself.com, 2016; see Vallor, 2016: 195-202). Of course, self-tracking is not akin to Socratic self-examination. As Shannon Vallor writes, “a dataset is not a life” (Vallor, 2016: 202).¹⁸ Still, dedicated self-trackers believe in their own kind of self-examined life, and also in a *collectively*-examined life. This means users are driven to share copious quantities of data about themselves, often with a large and unspecified audience

¹⁸ As a virtue ethicist, Vallor asks: is the quantified self movement merely “the next phase of humanity’s historical quest for the examined life? Is it an entirely *new* vision of the good life and the path leading to it?” Her answer is no, because we prize the examined life not for the data it yields, but “for the transformative nature of the practice itself and the dignity it confers upon those who take it up” (Vallor, 2016: 196-202).

(Lanzing, 2016: 11). Historically, Fitbit tended to make its users' profiles and activity public by default at the website Fitbit.com, which explains how in 2011 approximately 200 users inadvertently shared details of their sexual activity (Hill, 2011). In response, Fitbit changed its default settings to private. Nonetheless, the culture of self-tracking continues actively to encourage broad disclosure, aided by devices so small, light and waterproof as to be barely noticeable (Lanzing, 2016: 12-13). Data is shared with other device users; but it can also be shared unknowingly with health insurance firms, employers and government institutions who encourage "pushed self-tracking" (Lupton, 2014: 7). Again, the challenges to privacy are increased when data from several sources is combined, such as when data from Apple's HealthKit is combined with data from a Fitbit. For Apple, this is a selling point: "When your health and fitness apps work together, they become more powerful. And you might, too" (quoted in Lanzing, 2016: 13). Insurance companies in particular are highly active in the self-quantification space (Patterson, 2013: 10-11).

The emblem of the company-based threat to privacy is arguably the "cookie", more accurately known as a "persistent identifier". This is a small parcel of data sent from a website and stored in a user's web browser. Between visits, websites forget who users are; but cookies enable each user to be identified by the internet's equivalent of a name-tag, in the form of "I'm customer #582091". Every time a user loads Facebook, say, the browser sends the Facebook cookie (or cookies) back to the server to notify the website of the user's previous activity (Schneier, 2015: 47). Originally, cookies were devised to make using the web easier, but they now facilitate the tracking of users, and they are proliferating. If you visited dictionary.com in 2010, the site installed more than 200 tracking cookies on your browser (Schneier, 2015: 48). Over time, more sophisticated variations have enabled companies to collect "clickstream" data, showing precisely where and when users click. In 2007, it was revealed that US Internet Service Providers (ISPs) including Verizon, Comcast and AOL were monitoring user clickstream data and linking it with identifiable customer records (Nissenbaum, 2010: 29). Meanwhile, the term "data exhaust" has emerged to describe the digital trail users leave behind in their online interactions: where they click; how long they stay on a page; where their mouse-cursor hovers; and more (Mayer-Schönberger and Cukier, 2013: 113). The result of all this comprehensive,

surreptitious tracking has been described as creating a “soul in the machine” that accurately captures the user’s tastes and preferences which can be packaged into a coherent profile and then sold (Nissenbaum, 2010: 29). This process has been systematised and streamlined, with companies such as Acxiom, Equifax and Experian charging steep fees for comprehensive dossiers of data on individuals (Mayer-Schönberger and Cukier, 2013: 100, 150). Such data usually changes hands without the user realising. For instance, many common smartphone apps share personal information with third parties without notifying the user (Zang et al., 2015: 2-3).

Thanks to cookies and other tracking technologies, personalised advertising is a growth industry worth billions of dollars annually (Malheiros et al., 2012: 579). Ranging from mass customisation to one-to-one marketing, personalised advertising comprises customised promotional messages based on information such as name, past buying history and demographic (Baek and Morimoto, 2012: 60). Advertisers are getting better and better at targeting individuals (Dwyer, 2015b: 127). Indeed, such advertising has become sophisticated and automated, harnessing artificial intelligence, fuzzy logic, decision trees and more to find hidden trends, patterns and relationships (Guo and Zhang, 2015: 24). Famously, Target created an algorithm to determine when a customer was pregnant. Their system was triggered by subtle cues, including the purchase of zinc, lotion and a handbag large enough for nappies. Then, rather than sending women personalised ads for pacifiers and maternity wear, Target tactfully inserted baby-related notices among other advertisements. “As long as we don’t spook her, it works,” said a Target executive (Duhigg, 2012: 14). Studies confirm that such advertising works; further, it can even lead consumers to change their self-perceptions (Summers et al., 2016). As personalised advertising becomes both more accurate and more common, the repercussions for privacy are clear: to make effective personalised ads, advertisers need as much personal information as possible. And the trend, increasingly, is that companies are selling *users* to advertisers. From a user point of view, the rate of change is remarkable. By March 2017, four years after one-line text ads appeared atop my inbox, I instantly begin receiving ads for products I have just searched for, such as running shoes or a gift for my daughter. On current trends, the highly individualised ads of *Minority Report* sound not just possible, but imminent.

Physical location, I suggested in chapter one, is the new organising logic of the web (Gordon and de Souza e Silva, 2011: 7). This has led to the emergence of phrases including “locative media” and “geospatial web” to cover practices such as geotagging, online mapping and the location-based capabilities of social media (Crawford and Goggin, 2009; Dwyer, 2015b). In 2009, Twitter announced it would incorporate location data into tweets; in 2010, Facebook attached location data to its status updates; that same year, Google started to incorporate location data into every search, whether via the location settings on a mobile or the IP address of a desktop (Gordon and de Souza e Silva, 2011: 9). The dominant metaphor of the web has shifted from virtuality to mobility, and users’ location within the net they increasingly inhabit is being used to identify, quantify and understand them. Your mobile phone carrier knows where you are and, given the same knowledge about other users, knows who is with you. In chapter one, I noted that a 2012 study used location data to predict with great accuracy where users would be 24 hours later (Schneier, 2015: 1-2). This location data is then on-sold to a long list of companies (Almuhimedi et al., 2015). In 2012, researchers revealed that a single flashlight app was secretly collecting and on-selling location data from all its 50 million customers, including kids (Schneier, 2015: 46). Privacy concerns multiply when geolocational data is combined, say, with personal data and then shared with social media (Dwyer, 2015b: 125).

Further concerns arise from cloud computing, in the form of Apple’s iCloud, Google Drive, Microsoft OneDrive (formerly SkyDrive), Dropbox, Amazon Web Services and more. There is a fundamental shift, it seems, when I no longer store my documents exclusively on my own hard drive or my own smartphone. Certainly, the cloud is not invulnerable to hacks (Guthrie, 2016). However, my concern here is the way in which the movement of data out of devices *owned by users* and into a cloud *owned by companies* has ethical and legal ramifications about the rights to that data. Certainly, the cloud is growing at “a torrid pace” (McKendrick, 2016). Already, my phone regularly prompts me to backup all its contents to the HTC Cloud; I collaborate on documents stored at Google Drive and OneDrive; and friends share photo albums via Dropbox. Who owns this data? If I keep a copy of my thesis at Google Docs, can Google mine it for advertising? Use it for marketing? Sell it? The shift to cloud computing, it turns out, is not

simply symbolic. Across such services, users tend to retain ownership and copyright, but cloud computing services reserve the right, as specified in their terms and conditions, to access and use such data. As Google UK specifies:

When you upload, submit, store, send or receive content to or through our Services, you give Google (and those we work with) a worldwide licence to use, host, store, reproduce, modify, create derivative works (such as those resulting from translations, adaptations or other changes that we make so that your content works better with our Services), communicate, publish, publicly perform, publicly display and distribute such content (GoogleUK, 2014; see also Dropbox, 2015; and iCloud, 2015).

Often, a life lived in the net is simultaneously a life lived in the cloud. We check traffic on Google Maps, schedule our lives on Google Now and tag our kids in Google Photos. More likely, Google Photos automatically tags our kids for us. By design, these cloud-based apps work together. We pay our mortgage online; we talk to our lawyer via email; we transfer money. Thanks to big data and the internet of things, a more complete account of our lives is finding its way into the cloud, where it is stored in electronic facilities owned by private companies and scattered around the globe. Accordingly, the challenge to privacy grows. Privacy, as we shall see in chapter three, often involves control, and as our data migrates to the cloud, we lose a degree of control. Companies can and do delete our accounts, turn our data over to law enforcement authorities and store our data in countries where privacy laws are more lax (Schneier, 2015: 59).

Companies arguably present the biggest challenge to privacy in the internet age (Schneier, 2015: 47). This challenge is, in fact, the accumulation of many challenges. Above, I described Target living up to its name, targeting women it had secretly identified as pregnant with personalised advertising. A prospective employer would, presumably, be highly interested to know whether or not a potential recruit is pregnant. In many countries, it is illegal for employers to discriminate on the basis of pregnancy; using the internet, an employer might research and discriminate secretly. Research has shown that US employers consult social media when hiring, and that they discriminate against Muslim applicants (Acquisti and Fong, 2015: 4). As we have seen, users might be revealing their religion without realising. On current trends, the internet-based challenge to privacy from companies will snowball into greater challenges. As the predictive potential of location data reveals, this challenge may involve companies forecasting what consumers will be doing in the future. In fact, it may even

involve companies *influencing* what consumers will be doing. In 2007, Google's then chief executive Eric Schmidt said: "The goal is to enable Google users to be able to ask the question such as 'What shall I do tomorrow?' and 'What job shall I take?'" (Daniel and Palmer, 2007). In 2010, he repeated his claim. "I actually think most people don't want Google to answer their questions, they want Google to tell them what they should be doing next" (Jenkins Jr, 2010).

iii. The threat from governments

Winston Smith lives in a surveillance state, and he doesn't like the way citizens are constantly watched, the way free thought and autonomy are curtailed, and the way that anyone who expresses a dissenting opinion is rounded up and "re-educated" into orthodoxy. Unfortunately, because the government is all but omniscient, Winston's objections are soon noted, and he inevitably finds himself in the hands of the Thought Police, who take him to Room 101, deep in the fortified concrete bunker that is the Ministry of Love. "You asked me once, what was in Room 101," says O'Brien, a man Winston had thought was his friend. "I told you that you knew the answer already. Everyone knows it. The thing that is in Room 101 is the worst thing in the world" (Orwell, 1949: 296). When the government knows everything, it also knows each citizen's deepest fear. In Room 101, Winston sees two rats in a cage shaped to fit over his face. Faced with his worst nightmare, Winston breaks, yielding any resistance and betraying his lover.

A third threat is from governments. When talk turns to internet privacy, it is often this threat that dominates discussion, in the form of metadata retention, blanket surveillance or security agencies with three-letter acronyms. It is in this context that the Panopticon of Bentham and Foucault is commonly invoked, as a metaphor for the surveillance states made possible by the internet (Reiman, 2004: 196; Meikle and Young, 2012: 132-133; Schneier, 2015: 32, 97). And for some, Winston Smith presents an eloquent argument that an omniscient government can maintain and abuse power perpetually by exploiting citizens' weaknesses to stifle even the faintest murmur of dissent. Winston Smith, the protagonist of *1984*, has many sympathisers, who draw a direct link between surveillance states and oppressive totalitarianism (see chapter four). What's more, today's governments know more than the government of Winston's Oceania ever could have (Lessig,

2006: 208; see below). In part, this is because all the encroachments upon privacy described above, from both individuals and companies, can potentially also be accessed by government. On the net, as we have seen, data can easily be joined with more data, a phenomenon dubbed “cybernation” (Etzioni, 2015: 1264). Mountains of data are only one click away from being shared, a point illustrated, ironically, by the ease with which Edward Snowden was able to expose the surveillance practices of intelligence agencies.

In 2013, Edward Snowden was working for the National Security Agency in Hawaii when he copied an unknown number of classified documents. There may have been more than a million (Joye). In May, after exchanging encrypted emails with journalists, Snowden flew to Hong Kong and began sharing the contents. *Inter alia*, the documents exposed global surveillance programs undertaken by the governments of the United States, the United Kingdom, Australia and elsewhere, who were revealed to be engaged in mass, indiscriminate surveillance of domestic and foreign citizens, at times with the help of telecommunications and technology companies (Greenwald, 2014). The focus of Snowden’s revelations was his own workplace, the National Security Agency, or NSA. Formed in 1952, the NSA is a part of the military originally established to gather foreign intelligence. With the end of the Cold War in the 1980s and 1990s, the NSA became more focused on defence and more open, but this changed again following the terrorist attacks of September 11, 2001. Waging war on terror, the US government resolved to know everything about everyone (Schneier, 2015: 63). As revealed in NSA slides copied by Snowden, the NSA's stated objective now became to “Collect it All”, “Process it All”, “Exploit it All”, “Partner it All”, “Sniff it All” and “Know it All” (Greenwald, 2014). As Schneier writes, “Traditional espionage pits government against government ... But the terrorist enemy is different ... [its] members could be anywhere. Modern government surveillance monitors everyone, domestic and international alike” (Schneier, 2015: 63).

Meanwhile, the internet was transforming the nature of surveillance. Previously, a Chinese military network carried Chinese communications, a Russian military network carried Russian communications, and so on. The internet mixed everything together. If the NSA wanted to track terrorist emails, it would thereby track the emails of mums and dads too (Schneier, 2015: 64). The NSA doesn’t

limit itself to the internet: after September 11, 2001, the US government convinced the major telecommunications companies to hand over records to compile a database of every call ever made. Snowden's first published revelation was a document in which the FBI ordered Verizon to hand over the calling metadata of all its customers to the NSA (Schneier, 2015: 67). However, considerable surveillance *does* occur on the net, including three distinct NSA programs designed to collect Gmail data (Schneier, 2015: 63). And, as early as 2006, it was reported that the US Department of Justice had obtained search query records and other data from Google, AOL, Yahoo! and MSN (Nissenbaum, 2010: 30). The NSA also uses targeted programs. For "QUANTUMHAND", the NSA uses malware to disguise itself as a fake Facebook server to gain access to users' computers. When the user logs in, the NSA sends data packets that trick the computer into thinking they originate from the real Facebook. The NSA can then siphon out data from the computer's hard drive (Gallagher and Greenwald, 2014: 9). QUANTUMHAND is run by the NSA's Tailored Access Operations group (TAO), which has the job of hacking into computers remotely (Schneier, 2015: 71-72). As one malware program among many, QUANTUMHAND is integrated within the NSA's automated TURBINE system, which employs, according to TAO, "industrial-scale exploitation" (Gallagher and Greenwald, 2014: 3). NSA documents reveal plans to deploy "potentially millions of implants" as part of its "Owning the Net" program, which includes expanding TURBINE to enable "greater automation of computer network exploitation" (Gallagher and Greenwald, 2014: 4). There have been further revelations since Snowden. In 2017, WikiLeaks released documents revealing CIA surveillance programs. They included "Weeping Angel", which uses spyware to place a target's smart TV in "fake off" mode, in which a TV surreptitiously records conversations and sends them to a covert CIA server via the internet (Mitchelson, 2017b). Televisions thus become a sneakier version of Orwell's telescreens. There are two conclusions to draw: first, the internet makes it easier to collect data, as shown by QUANTUMHAND and Weeping Angel; and second, the internet enables all data that has been collected, from both online and offline sources, to be combined.

NSA surveillance is both general and targeted, and it is extensive. The NSA collects emails, text messages, browsing history, address books, location information and much more, although it is difficult to ascertain which data is

anonymised, which is analysed and which is retained (Schneier, 2015: 64-66). What's more, the NSA is just one of 17 US intelligence agencies; given that the NSA's existence remained secret for 20 years, there is a possibility that an 18th secret agency now exists (Schneier, 2015: 67). Surveillance also occurs outside these 17 agencies, including by the Bureau of Alcohol, Tobacco, and Firearms, which is building a large database to track people and their friends (Schneier, 2015: 69). Meanwhile, "fusion centers" have been set up for state and local law enforcement to gain access to data from national agencies such as the FBI. Initially set up to combat terrorism, these centres are now used for broader law enforcement, and have been used to spy on political protesters (Schneier, 2015: 69). Much of this surveillance is ill-defined and cloaked in secrecy. For Snowden, the result is "a system whose reach is unlimited but whose safeguards are not" (Greenberg, 2014).

Other countries' domestic surveillance programs are also extensive. NSA-like roles are played by the Government Communications Headquarters (GCHQ) in the UK and by the Australian Signals Directorate (ASD) in Australia. Similar organisations exist in Germany, France, Denmark, New Zealand, Israel, Canada, and elsewhere (Schneier, 2015: 70). In Australia, metadata retention laws were passed in 2015 granting approved government agencies warrantless access to two years' worth of customer call records, location information, IP addresses, billing information and other data stored by telecommunications companies. It is unclear which government agencies can access this metadata, as the names of some of the agencies seeking access have been suppressed, despite Freedom of Information applications (Duckett, 2016). In Russia, the System for Operative Investigative Measures is built into the internet, and is used against criminals, and also against journalists, human rights activists and political opponents; in China, more than 30,000 specialised police monitor the internet for phrases such as "Tiananmen" and "Amnesty International"; and in Thailand, India and Malaysia, people are regularly arrested based on internet conversations (Schneier, 2015: 70-71). It is impossible to know how many countries employ programs such as QUANTUMHAND, but it's fair to assume governments globally have access to sophisticated malware to collect emails, texts, call history, address books, to search history data and keystrokes, and to take screenshots, record audio, snap photos and monitor coordinates, before secretly sending back this information.

“It’s a reasonable assumption that most countries have these hacking capabilities. Who they use them against, and what legal rules control that use, depends on the country” (Schneier, 2015: 73-74). In 2016, evidence emerged of the Mexican government using spyware against anti-obesity campaigners, who began receiving highly personal text messages about friends’ funerals, unfaithful spouses and family members having serious accidents. One text said, “Simon buddy my dad just died we are devastated, I’m sending you info about the wake, I hope you can come”, with an attached link. The link contained an invasive spyware developed by NSO group, an Israeli cyberarms dealer. Like Italy’s Hacking Team and Britain’s Gamma Group, NSO Group claims to deal only with governments, and Mexico is a repeat customer. NSO spyware has been found on the phone of a human rights activist in the United Arab Emirates (Perloth, 2017).

Granted, government surveillance is not new. In eighteenth and nineteenth century France, the staff of the *cabinet noir*, or “black chamber”, worked behind the General Post Office in Paris, opening letters, reading their contents, then sealing them again without arousing the suspicion of sender or receiver (Cooke, 2013). In the twentieth century, the Stasi of East Germany became the paradigm of domestic surveillance by eavesdropping via hidden listening devices and encouraging children to spy on their parents (Mayer-Schönberger and Cukier, 2013: 150). However, the *cabinet noir* and the Stasi have been eclipsed by the extent of government surveillance in the internet age. In 2014, the Intelligence Community Comprehensive National Cybersecurity Initiative Data Center, or Utah Data Center, was opened. It has been described as the world’s largest black chamber, and has a storage capacity estimated to exceed a yottabyte. That’s enough to store everything ever written, plus every communication predicted to be made in the next century. Australia has been building a similar facility at HMAS Harman near Canberra (Cooke, 2013).

There are concerns beyond secret surveillance. In some cases, for instance, government departments anonymise citizens’ data. This is the case with the census, for instance, both in Australia and the US. However, researchers have shown that census data can easily be de-anonymised (Sweeney, 2000). At an international level, the proliferation of ePassports, which contain biometric data including a facial image on an RFID (Radio Frequency Identification)-enabled

chip, has raised issues (Juels et al., 2005). Then there is the growth of personalisation. “The dynamics of personalization shift power into the hands of a few major corporate actors. And this consolidation of huge masses of data offers governments (even democratic ones) more potential power than ever” (Pariser, 2011: 145). As we have seen, it is the combination of government and corporate watching that can be especially intrusive. In the context of a convergent, ubiquitous, multi-directional internet, government and companies comprise “a public-private surveillance partnership that spans the world” (Schneier, 2015: 78). The Snowden documents revealed the extent of the NSA’s reliance on telcos, search engines, software giants and other companies to collect data. Through programs such as PRISM, the NSA enlisted Microsoft, Google, Apple and Yahoo to provide information on specific individuals. Sometimes companies collaborate willingly; sometimes companies are compelled to comply; and sometimes the NSA and analogous agencies hack into company information without authorisation (Schneier, 2015: 78). If you run a business in the US, writes Schneier, and the NSA or FBI want to turn it into a mass surveillance tool, they can force you to comply, and then force you to keep the secret (Schneier, 2015: 84).

A fascinating side effect has emerged: in many countries, people who avoid the internet also stick out. These non-users become conspicuous for their very inconspicuousness. Often, however, such individuals may not be nearly as inconspicuous as they think. Thanks to shadow profiles and collation of offline data, such non-users may already have online profiles. And in the future, in the interests of security, some governments may mandate that every citizen has an online profile, thereby outlawing “hidden people” (Schmidt and Cohen, 2013: 33).

What, then, of predictive policing? In 1971, Germany began using computers and known data about family, housing, property, social situation and more to research the causes of criminality. This was intended to be the basis for more preventative police work (Goos et al., 2015: 63-64). In a phrase prefiguring the NSA’s “collect it all” mantra, the aim was for “everyone to know everything”. By 1979, the Federal Criminal Police Office, or BKA, had registered the names of 4.7 million persons, had fingerprints of 2.1 million suspects and photos of nearly as many. However, the system proved frustratingly ineffective, leading German courts to

curtail the use of computer databases for dragnet operations, let alone the prosecution of future crimes (Goos et al., 2015: 63-64). More recently, however, the internet has resuscitated the idea of predictive policing. In the US, half of all states use big data to help predict recidivism, and hence to determine whether an individual should be released or kept in jail; further, several precincts use big data to decide which streets, groups and individuals should be subject to extra policing; and, since 2006, police in Memphis have used the Blue CRUSH program (Crime Reduction Utilizing Statistical History) to target particular locales at specific times (Mayer-Schönberger and Cukier, 2013: 158). In the US case of *Wisconsin v Loomis*, a defendant was given a lengthy custodial sentence based in part on an algorithm devised by a private company that deemed him “high risk”. The defendant challenged his sentence because he had not been allowed to see the algorithm’s workings, but an appeal court found against him, holding that mere knowledge of the algorithm allowed for sufficient transparency (Tashea, 2017). As Mayer-Schoenberger and Cukier write, “The unsettling future *Minority Report* portrays is one that unchecked big-data analysis threatens to bring about, in which judgments of culpability are based on individualized predictions of future behavior” (Mayer-Schönberger and Cukier, 2013: 158). Granted, arresting an individual for a crime they are yet to commit remains the stuff of fiction. However, surveillance and big data are ensuring that predictive algorithms play an increasing role in modern policing (see Mayer-Schönberger and Cukier, 2013: 157-163). Unfortunately, predictive algorithms contain all the biases, blindspots and flawed assumptions of those who engineered them (Vallor, 2016: 193). The ensuing challenge to privacy is more dramatic than a Spielberg blockbuster.

Currently, the world’s most extensive, effective spying network belongs to the US. Naturally, other countries want to harness its knowledge. Meanwhile, the US wants data collected by other countries. Hence the collaboration between the countries of the “Five Eyes”: the US, the UK, Australia, New Zealand, Canada. And of the Nine Eyes, which also includes Denmark, France, the Netherlands and Norway. And the Fourteen Eyes, which adds Germany, Belgium, Italy, Spain and Sweden. What’s more, the NSA also works with India, Saudi Arabia, and other countries, on top of working especially closely with Israel. Schneier writes:

All of this gives the NSA access to almost everything ... The endgame of this isn’t pretty: it’s a global surveillance network where all countries collude to surveil everyone on the

entire planet. It'll probably not happen for a while ... [but] it's the rational thing to do (Schneier, 2015: 76-77).

This surveillance will potentially become even more significant as governments start to “nudge” citizens into certain types of behaviour. Co-owned by the UK government, the Behavioural Insights Team has had success using behavioural psychology to influence people via subtle cues into paying fines on time and consuming fewer fizzy drinks. In Australia, the group has worked with the New South Wales government to keep commuters out of the CBD at peak times and to reduce domestic violence (Miller, 2016). This is potentially problematic, given that human agency requires that our moral practice is not passive, but “our own conscious activity and achievement” (Vallor, 2016: 203). These are pro-social initiatives; presumably the same techniques could be harnessed for anti-democratic aims, such as manipulating the way citizens vote. Research has already shown how Facebook can manipulate the outcome of elections (Brand, 2016). The potential for governments and companies to collaborate on exploiting private information to keep incumbents in power already exists.

iv. Resistance

The picture I have painted thus far suggests everyone can know everything about everyone. It suggests that a click or two will reveal all about someone to individuals, to companies and to governments. However, this is patently false. The theory of Panopticon 2.0 does not match the reality. There are limits to omniscience, even for those at the top of the surveillance pyramid, such as Acxiom or the NSA. For instance, there remain a number of governments who do not share their information with the US, such as Russia and Iran, and who are likely to continue on their own path (Schneier, 2015: 77). Moreover, there are instances of corporate pushback, with Google, Facebook and Yahoo! all claiming to have resisted requests to give up data to US intelligence agencies (Schmidt and Cohen, 2013: 267). In 2016, the FBI revealed it had obtained the iPhone 5C of a terrorist who had killed and been killed in San Bernardino, California. After the NSA was unable to break into the operating system, Apple refused an FBI request and ensuing court orders to un-encrypt the phone, arguing that creating a

backdoor would put the security of all its users at risk (Levy, 2016b).¹⁹

Companies and governments do not always collaborate. What's more, individuals usually have limited means at their disposal. Some privacy does persist, even on the net.

After all, not everyone is on the net. By the end of 2016, 53 per cent of the world's population, or 3.9 billion of the world's 7.4 billion people, remained disconnected. They were, overwhelmingly, from the world's poorest nations (UN, 2016: 6). For the 47 per cent who are internet-connected, privacy-enhancing strategies do exist. One chief strategy is anonymity (Meikle and Young, 2012: 127-147; Coleman, 2014). Initially, all web use was presumed to be anonymous. Since at least 2000, however, it has been recognised that net anonymity is under threat, and must be achieved through effort (Sobel, 2000: §22). Hackers in particular have a rich history of working under the cloak of anonymity, best typified by the activist collective Anonymous (Coleman, 2014). Another common use of digital anonymity is in the medical profession, where individuals are given pseudonyms that enable their medical history, but not their identity, to follow them throughout the system (eg, Elgesem, 1996). In Australia, users can register for a pseudonym for their "My Health Record" (MyHealth, 2016). Entire corners of the web are devoted to anonymous interactions, including the bulletin board 4chan, where users post content that may be embarrassing, shocking or illegal (Meikle and Young, 2012: 144-145). Among other subjects, boards exist that are devoted to "anime & manga", "technology", "weapons", "fashion", "LGBT", "hardcore" and "adult GIF" (4chan, 2016). It is worth noting, however, that 4chan has been associated with misogyny, revealing how privacy can cloak sexism (Jane, 2016a: 30-31; see chapter three).

By using a browser such as TOR, users can anonymously navigate the net, including the deep net (see chapter one). Some parts of the net are accessible only via anonymous browsers such as TOR, which works by distributing a user's activity over several places on the net, so that she can't be linked to one particular place (Rudesill et al., 2015). The ensuing anonymity can be used both for pro-social and anti-social purposes. Above, I discussed the way wearables can

¹⁹ The standoff ended when the Department of Justice announced it had broken into the iPhone with the help of hackers paid by the FBI (Perlroth, 2016). As such, this is hardly an uncomplicated example to show that there are indeed limits on the free flow of data.

encourage users to be fit and healthy; conversely, pro-anorexia communities, known as “pro ana” and “thinspo”, also thrive online. Offline, these communities struggle to coalesce; on the net, shielded from social norms, anonymity emboldens thinspo teens to meet (Arseniev-Koehler et al., 2016). Always, however, the question arises: just how anonymous is anonymous? Not very, writes Schneier: “Maintaining Internet anonymity against a ubiquitous surveillor is nearly impossible ... Anonymity is fragile. We either need to develop more robust techniques for preserving anonymity, or give up on the idea entirely” (Schneier, 2015: 42-45). However, anonymity potentially *can* prevail. The inherent design of the internet means that we can’t attach identifiers to data packets on the net, and we can’t verify the identity of someone somewhere sitting in front of a computer (Schneier, 2015: 132). Strong anonymity is only vulnerable in the face of mass surveillance and persistent effort. On the net, anonymity has the potential to be a powerful, if not infallible, privacy-enhancer.

The theory of Panopticon 2.0 suggests that all can see all. However, I can’t simply log into a government agency’s files to see what they contain. Further, I can’t simply access all the data Facebook or Google has, just as I wouldn’t know how to hack into my neighbour’s hard drive. Security measures are in place to stop me gaining such access. For Derek Bambauer, security implements privacy (Bambauer, 2013: 671). Encryption can be particularly effective. One version is SSL, or Secure Sockets Layer, which builds an encrypted link between a web server and a browser. Google is working towards using SSL encryption as a standard (Yorke); Apple relies heavily on encryption (Levy, 2016b); and in 2016, Facebook adopted end-to-end encryption for its Messenger instant messaging service (Yadron, 2016). Potentially, encryption is a tool of great power (Schneier, 2013: 143-144). As Edward Snowden says: “Encryption works. Properly implemented strong crypto systems are one of the few things that you can rely on” (in Greenwald, 2013). Techniques such as encryption (to which I return in chapter six) show that checks and blocks exist.

In other ways, too, users are adopting tactics to preserve their privacy. One recurrent strategy is obfuscation, the deliberate use of ambiguous, confusing and misleading information to thwart the collection of data by hackers, companies and governments (Nissenbaum and Brunton, 2015). Obfuscation is intended to muddy

the waters and make it impossible for users to be accurately profiled; methods include evasion, noncompliance, refusal and sabotage. In this way, it is argued, average users are also signaling that they want to use the web but do not want to be tracked and profiled. One obfuscatory tool is TrackMeNot, a free browser extension that dupes trackers into recording false user activity as well as real user activity. It works:

... not by means of concealment or encryption (i.e. covering one's tracks), but instead, paradoxically, by the opposite strategy: noise and obfuscation. With TrackMeNot, actual web searches, lost in a cloud of false leads, are essentially hidden in plain view (TrackMeNot, 2016).

Another free browser extension is Ghostery, which enables users to identify and block cookies and other tracking devices, while DuckDuckGo is an American search engine that doesn't track its user's online activity (Lymn, 2012).

Disconnect.me is software developed by former Google employees which can block trackers and search privately, and offers stronger protections for a fee: "Our privacy policy, in a sentence: we don't collect your IP address or any other personal info, except the info you volunteer" (Disconnect, 2016). Meanwhile, the data-hoovering tendency of Facebook has prompted the development of alternative social networks, such as Diaspora, which is not-for-profit and user-owned (Christians et al., 2012: 99). Another alternative is Snapchat, an instant-messaging service on which messages self-destruct, which has become popular among teenagers and young adults, often for exchanging selfies and humorous content (Piwek and Joinson, 2016). Many and varied, strategies of resistance can bolster privacy in the face of surveillance, sousveillance and lateral viewing.

Research is showing that people care deeply about digital privacy, even if their actions don't always align with that care (Debatin et al., 2009; Taddicken, 2014). Often, however, their actions do align with that care. Anonymity, encryption and other privacy-protecting strategies confirm it. So does a 2016 study showing the security and privacy concerns have stopped many users posting to social networks, expressing opinions or buying from websites. For these users, the internet can no longer be trusted for everyday activities (Peterson, 2016). In 2016, a widely circulated photo showed Mark Zuckerberg at his desk, where a piece of stickytape had made the webcam of his laptop inoperable (Titcomb, 2016). As we

saw in chapter one, Zuckerberg says privacy norms have shifted and loosened; clearly, he still wants some.

II –“LO”, the Internet! An epic clash of norms and values

Some people have told me that internet ethics is a contradiction in terms. I closed the previous section, however, with research showing that people still value privacy, even on the internet. In this section, I explore some of the ethical values and norms that prevail on the internet. First, I address the widespread perception that the internet is a place of lawlessness, a virtual wild west where, if you cross the wrong person, you’ll wind up hacked to pieces. For many, I argue, the net is not just a locale where conventional ethics have no place, but a locale where they *should* have no place. However, this view is misguided. As I proceed to argue, ethical standards already do apply on the net. Elaborating these standards is tricky, however, because they are complex, contradictory and sometimes unclear. This is a point I continue to prosecute in the second part of this section, where I argue that certain values and norms have been embedded in the architecture of the net and its platforms. These embedded values and norms have changed over time and are sometimes conflicting, but, generally, privacy has not been one of them. Finally, in the third part, I argue that norms should apply on the net, including for matters of privacy. The question becomes: which norms?

i. User-generated norms

The internet is at once a continuation of the media that preceded it, and a new domain where individuals, companies and governments are staking claims, competing to impose their own values, norms and regulations. It is, in short, a place of “complex relations of contestation and continuity” (Meikle and Young, 2012: 9).

In extremis, the ethical contest manifests as a battle of utopians against dystopians (Dinello, 2005; boyd, 2014: 15, 24). The former hold that the internet is so fundamentally good that it is humanity’s saviour. Google famously hinted in this

direction with its slogan, “Don’t be evil” (Morozov, 2011).²⁰ According to “techno-utopians” and “cyber-optimists”, the internet is fated to promote knowledge, democracy and egalitarianism, thereby liberating humanity. On this view, the internet enables the unfettered flow of information, inspires creativity and innovation and promotes equity, democracy and justice. “Never before in history have so many people, from so many places, had so much power at their fingertips,” write Google’s Schmidt and Cohen (Schmidt and Cohen, 2013: 3-4; see also Kurzweil, 2005). For detractors, however, the internet is an unmitigated force for evil, shaping up as the final nail in humanity’s coffin. Inspired by books such as *1984* and films such as *The Matrix*, the “techno-dystopians” and “cyber-pessimists” contend that the internet and associated technologies are precipitating the end of civilisation as we know it. On the net, they say, we are superficially connected but profoundly isolated (Turkle, 2011). In 2014, cosmologist Stephen Hawking said that artificial intelligence “could spell the end of the human race” (Vallor, 2016: 1). Indeed, the perception that society is under threat from the internet and other digital technology has become so widespread that it has spawned the terms “technophobia” and “technopanic” (Dinello, 2005; Marwick, 2008). As machines take over mankind, it is argued, the internet will enslave us.

In recent years, the volume of utopian and dystopian rhetoric has quietened. A more enduring claim, however, is that the net is a fitting home for anarchy. As Schmidt and Cohen write: “The Internet is the largest experiment involving anarchy in history. Hundreds of millions of people are, each minute, creating and consuming an untold amount of digital content in an online world that is not truly bound by terrestrial laws” (Schmidt and Cohen, 2013: 3). There are at least three reasons for the spread of this anarchist spirit, which pervades internet groups such as WikiLeaks (Curran and Gibson, 2013). First, it stems from the net’s inherent complexity. In a constant state of flux, the net is ever changing, an exercise not just in interactivity, but in intercreativity. The most defining characteristic of the internet may well be its very lack of definition (Bane and Milheim, 1995: 32). Castells describes the internet as a “space of flows”; it is relentlessly protean, stubbornly mercurial as its moment-to-moment reinvention continues (Castells, 1996: 376-378). This can lead to ethical inertia. As Vallor notes, when our future

²⁰ Initially, Google’s “Don’t Be Evil” slogan was invoked to signify the search engine’s strategy of not accepting advertising; only later did it come to be invoked more broadly as a company-defining ethic. Once Google started accepting ads, the motto was phased out (Morozov, 2011).

is opaque, it is hard to be motivated by an ethical principle (Vallor, 2016: 6). There is thus an opening for a spirit of anarchy. Second, the internet enables us to interact in unprecedented ways, making it difficult to deduce which ethical principles ought to apply in particular instances. If I type “how to make a bomb” into Google, is the government justified in collecting and keeping a record of that search, and then placing me under intensive surveillance? If I see a work colleague on dating app Tinder, is it ok to take a screenshot and email it to all my friends, some of whom are mutual colleagues? And if I take an embarrassing photo of a drunken bacchanal, can I upload the image to Facebook and tag all my friends who appear in the image? Users might ask themselves: if it is unclear which norms should apply, then why should I not be able to apply whichever norms I like? And third, the net is sometimes described as a rightful home for anarchists simply because many users believe that conventional ethics do not and ought not apply on the net. They believe that the net is a new hope for humankind, which offers all the promise of a Hobbesian state of nature, or at least a libertarian paradise, where freedoms are unfettered by the state, and individuals’ interactions can be unhindered by laws and leaders (Streeter, 1999: 50). On this view, the net is a virtual wild west. This frontier ideology is made explicit among hackers, where good guys are known as “white hats” and bad guys as “black hats”. The former create and build, the latter destroy, extort and wield malicious intent (Zetter, 2016; Meikle and Young, 2012: 30-31; see below). The implication is that in the untamed frontier of the internet, men (and occasionally, perhaps, women) must take the law into their own hands.

The anarchic aspirations of some net users is evident in the way that DIY justice proliferates online, including under the banner of “electronic civil disobedience”, or ECD (Meikle and Young, 2012: 142-144). One notable example is the hacker group Anonymous, which engages in co-ordinated activism, or “hacktivism”, and whose ideals have been described as “cyberlibertarian” (Goode, 2015).

Anonymous have hacked in support of WikiLeaks, the Occupy movement and Arab Spring protesters, while their targets have included corporations (Mastercard, VISA, Paypal), religious groups (the Church of Scientology, Westboro Baptist Church) and terrorists (ISIS). Their methods include Distributed Denial of Service (DDoS) attacks, which can overwhelm targets’ websites and render them useless, and the group’s emblem features a man without a head,

suggesting its leaderless structure and anarchic ideals (Coleman, 2014).

Sometimes vigilantes organise themselves into *ad hoc* gangs. In the days after the Boston Marathon bombings of 2013, users gathered in forums to crack the identity of the bomber. Within days, they identified missing university student Sunil Tripathi, launching a digital lynch mob on Tripathi's family, who felt under siege. The siege lasted overnight, until the real bombers were identified (Shih, 2013). Another brand of digital vigilantism, or digilantism, is that of feminists who seek to call out misogynistic behaviour, including rape threats. Groups that engage in feminist digilantism include Destroy the Joint and Sexual Violence Won't Be Silenced, or SVWBS (Jane, 2016b). Online, vigilante behaviour is regularly linked with compromising privacy, involving a "weaponised visibility" which enables a parallel form of criminal justice and policing (Trottier, 2016). "Digital vigilantism is a process where citizens are collectively offended by other citizen activity, and coordinate retaliation on mobile devices and social platforms." Such retaliatory activity can include doxing, as described above, and tends to be unwanted, intense and enduring (Trottier, 2016: 1).

In 2016, images of Australian schoolgirls were shared to a pornographic site. When one girl asked for photos to be removed, she received the response: "Darling, don't be a slut and you won't end up here. Once a photo is on snapchat or the Internet, it belongs to the Internet" (Olding and Munro, 2016). Online, conventional ethics are sometimes treated with disdain. On this account, the internet is another world, where the rules of the real world don't apply. Despite such remarks, however, the net is not an ethics-free zone. Rather, it is a place where norms, ethical and otherwise, are fiercely contested. To say that a revealing photo of an underage schoolgirl "belongs" to the internet is to make an ethical claim, albeit a claim that runs counter to conventional ethical norms. Similarly, some forms of digital vigilantism, including many of the activities of Anonymous, suggest that ethical norms should be applied online, but that these norms differ from offline norms. By contrast, other instances of digital vigilantism, such as the (mistaken) identification of Sunil Tripathi as the Boston bomber, suggest that offline ethics and justice *should* prevail online, but that netizens must take matters into their own hands because there is insufficient online enforcement of such standards. And so too a moral code applies to doxing, even if that moral code is inconsistent and sometimes problematic and misogynistic. In short, ethical norms

are being applied on the net, even if those norms are contradictory and sometimes dangerous.

In some cases, this is having extreme results, as with the normalisation of “hurtcore porn”, which involves real footage of people, including children, being sexually abused. As one hurtcore user says, “At first I felt ashamed in myself for being attracted to such a thing, but as time went on I slowly grew more accepting of myself” (Thomas, 2015). There is a similar risk of privacy norms shifting not because they should, but because people behave according to new norms they wouldn’t usually accept (for reasons such as peer pressure). These norms can then become entrenched after a period of habituation. Online, there exists “the risk of collateral privacy damage and the plasticity of norms” (Hull et al., 2011: 295). The internet is not ethics-free and it is no anarchists’ playground; it is neither libertarian paradise nor Hobbesian state of nature. Rather, its users are engaged in an ongoing clash of ethical norms.

ii. Embedded values

So far, I have been arguing that the internet’s user-generated norms are both a continuation and a contestation. Further, they are confusing and contradictory. These are the net’s *user-imposed* values; but what about the net’s *embedded* values?

The internet, in the form of ARPANET, flickered to life in October 1969, with an inaugural message that had been intended to read, “LOGIN”. Unfortunately, the receiving computer crashed, cutting the message to “LO” (Hafner and Lyon, 1996: 153). It was an Old Testament beginning for a new medium, just as the Cold War was at its iciest. Facing a hostile Soviet Union, the United States Defense Department Advanced Research Projects Agency (ARPA) wanted a communications network that would survive nuclear attack (Meikle and Young, 2012: 29). ARPANET was thus built to link thousands of autonomous computer networks in a seemingly limitless number of ways, thereby circumventing barriers and breakdowns (Castells, 1996: 6-7). To suit the military, the network was closed; but it was nonetheless intended to be the epitome of decentralisation. In

these early years, an ability to circumvent obstacles proved a defining characteristic (Castells, 1996: 19, 26).

When the Iron Curtain came down, the internet stayed up. However, its focus shifted. In the 1970s and '80s, the net was embraced by scientists and academics, who used it to access information stored on remote computers, and to enable distant colleagues to share their work (Leiner et al., 2009: 27; Hafner and Lyon, 1996: 240-244). By the time the Cold War thawed in the late '80s, the military's ARPANET had been superseded by the National Science Foundation's NSFNET as the internet's high-speed "backbone" (Bane and Milheim, 1995: 2). As scientists and academics began to unlock the new medium's collaborative potential, an ethic of openness was superimposed onto the net's architecture of decentralisation. On one view, the net's academic, peaceful origins are more significant than its military origins. ARPANET, write Hafner and Lyon, "embodied the most peaceful intentions to link computers at scientific laboratories across the country so that researchers might share computer resources. ARPANET and its progeny, the Internet, had nothing to do with supporting or surviving the war – never did" (Hafner and Lyon, 1996: 79-80). The original intention behind the net may well have been more peaceful than martial. Some of the internet's architects argue that any suggestion that the internet was intended to be a nuclear resistant network is merely "false rumor" (Leiner et al., 2009). In any case, an architecture of openness came to prevail; the pioneers of the net wanted "to allow the network to evolve as an open system of computer communication, able to reach out to the whole world" (Castells, 1996: 19).

During the 1970s and 1980s, the values of decentralisation and openness were joined by a third influence, the "hacker ethic". It was eccentric, anti-authoritarian and often libertarian (see Streeter, 1999: 49-50). By the mid-1970s, tens of thousands of the world's brightest microelectronics innovators were working in California's Silicon Valley (Castells, 1996: 55). These internet pioneers met in loose clubs, such as the Home Brew Computer Club, and their members included Bill Gates, Steve Jobs and Steve Wozniak, who between them created Microsoft, Apple, and 20 more companies. They blended innovation and invention with informality and irreverence and saw themselves, in a way, as successors to the counter-culture of Haight Ashbury. In the 1960s, San Francisco had been the

epicentre of the world's hippie counter-culture, complete with its persistent attempts at sexual, gender, racial and other emancipations; by the 1970s, 50 kilometres to the south, Silicon Valley became the epicentre of the world's nascent computer industry. Unsurprisingly, there was overlap in their outlook (Castells, 1996: 6, 54). "Hackers built the Internet," write Mickle and Young, and the original hacker ethic was built on principles including mistrust authority, promote decentralisation and keep all information free (Mickle and Young, 2012: 30; Levy, 1984: 26-36). The Hacker Ethic was codified by Steven Levy, with six precepts that included, "You can create art and beauty on a computer," and "Computers can change your life for the better" (Levy, 1984: 26-36). The idea was to be playful and probing, and thus to spark a "benevolent ripple" through society by which "computers would indeed change the world for the better" (Levy, 1984: 36). The hacker influence still persists, and has been coded into the net to protect freedom of information, freedom of speech and the principle of "net neutrality", which prescribes a free and open internet (Godwin, 2003). The hacking spirit is evident in open-source software, or OSS, and in the work of coder Richard Stallman, who campaigned for *free* software - software that can be used, studied, distributed and modified (Stallman, 2002).

As it evolved in the '70s and '80s, then, the internet was largely built with an architecture that fostered decentralisation, openness and hacking. Sometimes there was tension in these values. A comprehensive account of the rise of the net needs to allow for both the "closed world" of the Cold War and the open, decentralised world of the antiwar movement; only thus can we determine whether today's internet tends to foster democratisation or control (Rosenzweig, 2004).

Overwhelmingly, however, the values encoded into the net in its first two decades tended to promote the free flow of data. Openness, collaboration and sharing were encouraged; privacy was not. Just as the net's creators had not foreseen email or social media, neither had they foreseen the privacy concerns that would arise from email, social media and location-based apps, among other developments. On the net, privacy concerns were not initially addressed by design, and could only later be addressed via add-ons (Michener, 1999).

In the 1990s, as the internet went mainstream and global, corporate and government values started to be encoded into the internet and its platforms.

Following the launch of the world wide web in 1991, Internet Service Providers (ISPs) began to turn the net into a commodity (Leiner et al., 2009: 30). Soon, the net's free spirit was joined by a free enterprise spirit. First came the proliferation of commerce-based websites; then, in the 2000s, social media reinvented the web as "Web 2.0" (Meikle and Young, 2012: 65-68). Companies continue to play a central role. The reach of Facebook has already been addressed; and, in February 2016, Alphabet, the parent company of Google, became the world's largest publicly traded company, surpassing Apple (Levy, 2016a). Online as offline, a key motivator for companies is profit, and on the net profits are often realised at the expense of privacy, particularly for social networks (Elmer, 2013). Hence many companies embed an ethic of openness into their services, and particularly into their privacy settings. On Facebook, for instance, a user is automatically subject to behavioural advertising unless she opts out, and this targeting is based on information drawn from user activity both on and off Facebook (Facebook, 2016a). The requirement of opting out is part of a wider trend: on the internet, default settings tend to be public. If a user wants privacy - for instance, by avoiding targeted advertising - she generally needs to opt out of these public defaults (Spinello, 2011: 43-45; Christians et al., 2012: 97-101). In the face of these default settings, users are often apathetic. One study of 4000 US students found that few altered the default settings for privacy, which are highly permeable (Gross and Acquisti, 2005). Certainly, when users choose privacy settings, they often decide based not on reasoned choice, but on the settings of their peers (Lewis et al., 2008). The upshot is that users regularly regret some online disclosures (Sleeper et al., 2013; Wang et al., 2014). The design of Facebook, and many other platforms, invites confession, revelation, exposure and sharing. As Timothy Dwyer writes,

Powerful market-dominating new media corporations, such as Google (the owner of YouTube), Facebook, LinkedIn, and Twitter, have made it clear that it is their avowed intention to reconfigure people's understandings of the meanings of personal privacy (Dwyer, 2015b: 121).

Given this corporate drive for profit and publicity, writes Dwyer, "ethical standards are in an uphill battle for survival" (Dwyer, 2015b: 126).

In the '90s, just as companies were becoming active online, governments also started playing an increasing role. This role was both visible (in laws seeking to

regulate online behaviour) and hidden (in the covert surveillance described above). Hence it has been argued that the original internet was anarchic, but has now become a highly regulated and controlled space. “The invisible hand of cyberspace is building an architecture that is quite the opposite of its architecture at its birth,” writes Lawrence Lessig. The challenge now, Lessig argues, is not an excess but a shortage of freedom (Lessig, 2006: 4-5). This shift from openness to control is evident, for instance, in the prosecution of Aaron Swartz, a hacker who co-founded Reddit and Creative Commons. Part of the “internet free culture movement”, Swartz was arrested (and later committed suicide) after publishing and sharing academic articles otherwise available only for a fee (Peters, 2016: 1-3, 10). Initially, decentralisation, openness and the hacker ethic were coded into the net; these have now been joined, and sometimes overrun, by the values of companies and governments, which include profit and surveillance. Consistently, these values work against privacy.

In some ways, admittedly, privacy has found a place on the net. Here and there, privacy has also been encoded. I have already discussed anonymity, encryption and security. In the ‘90s, for instance, the success of the commerce-driven “Web 1.0” required an in-built level of security to protect users who were handing over credit card details. This security, when effective, brought a degree of privacy. As we have seen, strategies are emerging for those who wish to maintain their privacy. As the amorphous and unpredictable internet continues to grow and shift, there is a push-pull of competing imperatives. This is evident with encryption. Post-Snowden, tech companies have been implementing strong encryption by default, to thwart “thieves, vandals and foreign powers” (Levy, 2016b). Meanwhile, various governments have at the same time been working either to circumvent cryptography, or to outlaw it. In India, Russia, Saudi Arabia, the UAE and Indonesia, governments threatened to ban Blackberry if the company didn’t unencrypt its communications. Given that BlackBerry has not been banned, it seems fair to surmise that governments in those countries are now eavesdropping on BlackBerry exchanges (Schmidt and Cohen, 2013: 73; Schneier, 2015: 70-71). In any case, the NSA’s BULLRUN program and the GCHQ EDGEHILL program have both been successful against much of the security common on the internet (Schneier, 2015: 85). In 2016, the US government rebooted its war against encryption by enlisting a way to hack into Apple’s iPhone (Levy, 2016b). The

story is getting more complicated, but, generally, openness still tends to prevail. Privacy has tended to be encoded into the internet only recently and sporadically, via add-ons and plug-ins, and via specific privacy-promoting platforms, companies and strategies. Privacy can be had online. Often, though, it takes work. Data tends to be public by default, private by effort.

One final embedded value, popularity, is also worth considering. On the internet, popularity is privileged. This is most evident on social networking sites, where users are encouraged to have as many friends and followers as possible. On social media, popular users are perceived to be more socially and physically attractive, extroverted and approachable than unpopular users (Scott, 2014). Online, a user's worth can seemingly be measured by number of friends, followers, likes, shares, favourites, retweets and comments. For small sums, users can buy fake friends from sites offering "best quality friends on the market" (SocialYup, 2016). People have always connected, revealed and shared. The difference is that on the net, driven by the imperative of popularity, users are encouraged to do so with less discretion and discrimination. The architecture of social media, and the net generally, encourages liking, sharing and befriending. On the internet, challenges to the condition of privacy are often encoded in the form of embedded values, including popularity, and these challenges are significant, frequent and persistent.

iii. Net ethics: non-determined and necessary

I have been talking just now of the values and norms coded into the net and its platforms. One important proviso is that I do not mean to imply thereby that the internet will, to a greater or lesser extent, determine our future. In other words, I am not arguing for technological determinism, the notion that society and human life are determined by technology. In the hard determinists' vision of the future, "we will have technologized our ways to the point where, for better or worse, our technologies permit few alternatives to their inherent dictates" (Smith and Marx, 1994: xii). Intuitively, the hard determinist position seems false (Tomlinson, 2007: 11-12). Humans build technology; humans use technology. It is in the building and the usage, not in the technology itself, that outcomes would seem to be decided.

Rather, I am arguing that the net's prevailing ethics have two sources: the values and norms imposed by users; and the values and norms that are embedded, or encoded, in technology (Lessig, 2006; Spinello, 2011). Rather than technological determinism, I am arguing that people and technology have an effect on each other, just as people and media have an effect on each other, and hence as people and the internet have an effect on each other. In the words of sociologist Manuel Castells, there is a "dialectical interaction between society and technology": "technology does not determine society: it embodies it. But neither does society determine technological innovation: it uses it" (Castells, 1996: 5 fn.2). Or, as US historian Melvin Kranzberg wrote, in the first of his six succinct laws on technology: "Technology is neither good nor bad; nor is it neutral" (Kranzberg, 1986: 545). The internet too is neither good nor bad; nor is it neutral. Rather, there are powerful values and norms at play. Online, there is a continuation and contestation of norms, some of them user-imposed, some of them embedded in the very architecture of the net and its platforms.

In less than half a century, the internet has grown into a global, horizontal computer network accessed first by thousands, then millions, now billions of people. The contest of norms and values has grown too. On the net, user-generated ethical claims range from conventional to eccentric, from safe to toxic: the net is utopian; it is dystopian; it is and should be a place of anarchy; conventional justice must be meted out by digital enforcers and cyber lynch mobs. Meanwhile, norms and values have been coded into the internet and its platforms. As a result, the condition of privacy has diminished, and the right to privacy has become less clear. A surveilled openness has come to prevail, enabling hackers to phish in the cloud, Acxiom to build extensive dossiers and governments to monitor meticulously. As the net tends to convergence, ubiquity and multi-directionality, values and norms favouring openness are becoming yet more pervasive and dominant.

The internet, like preceding media, has been regularly demonised as a site of social disruption (Meikle and Young, 2012: 194). Certainly, the net is an ethical battleground. In some ways, admittedly, privacy has found a place on the net. As a result, each of us is not an actual Panoptes. Even companies and governments are not omniscient. Once we discard technological determinism and discount

anarchic aspirations, we can acknowledge that certain values and norms already prevail on the net. What's more, we can acknowledge that norms and values *should* prevail on the net. When it comes to privacy, all that remains is to decide what those norms and values ought to be, before then elaborating how best to protect them.

Conclusion

In the opening scene of *Minority Report*, a husband is in a jealous rage. Having caught his wife in bed with another man, the husband is about to stab them both with a pair of scissors. Just in time, a cop named John Anderton busts in and intervenes. "Mr. Marks," he says, "by mandate of the District of Columbia Precrime Division, I'm placing you under arrest for the future murder of Sarah Marks and Donald Dubin that was to take place today, April 22, at 0800 hours and four minutes." Mr Marks is stunned: "I didn't do anything! I wasn't gonna do anything!" By the end of the film, however, the Precrime department is shut down. As Anderton says in a final voiceover: "In 2054, the six-year Precrime experiment was abandoned. All prisoners were unconditionally pardoned and released, though police departments kept watch on many of them for years to come." In this denouement there is an explicit recognition that condemning people for future actions is fundamentally problematic. At the very least, the film suggests that the ethics of predictive technology must be carefully considered and thoroughly investigated.

I began this chapter by breaking down the challenge to privacy for internet users into three parts: the challenge from individuals; the challenge from companies; and the challenge from governments. A hacker who uses malware to access a webcam is unlike a company that creates a detailed profile of a non-user, which in turn is unlike a government seeking intelligence to thwart terrorism. These challenges are often overlapping, but nonetheless have distinct elements, and, to some extent, require varying responses. What's more, these challenges are not all-conquering. Restrictions on data flows exist. Anonymity is sometimes possible. Encryption can prevent surveillance. And a number of technologies and strategies have been developed to thwart surveillance and enhance privacy. In section two, I

proceeded to argue that the ethical norms that prevail on the internet are both a continuation and a contestation. For some users, the net offers all the promise of an anarchists' haven. On this view, the net should not be a place for conventional ethics. However, a complex of social and ethical norms *is* being applied online, including the norms that underpin vigilante justice and hurtcore porn. These user-generated norms are fiercely contested. Meanwhile, values and norms have been embedded in the architecture of the net and its platforms. These too are complicated and sometimes contradictory, but tend to promote openness, not privacy. These embedded values stem from: the decentralisation built into the net's original architecture; the collaborative and connective nature of academia's contribution; the commitment to freedom of information attending the hacker ethos of Silicon Valley in the '70s and '80s; the increasing corporatisation of the net, both in web 1.0 and then web 2.0, which encourages sharing; the burgeoning surveillance practices of governments; and the compelling imperative of popularity. People tend to care about privacy; the way they use the net tends to work against it.

The future of the internet is not determined. Rather, it will be written by how we code it, how we make it and how we use it. As several of the internet's architects wrote in 2000: "If the Internet stumbles, it will not be because we lack for technology, vision, or motivation. It will be because we cannot set a direction and march collectively into the future" (Leiner et al., 2009: 31). In order to *best* decide how to code, make and use the net, we need to chart an ethical course. This involves deciding, for instance, what sort of predictive policing and personalised advertising are fair and just. When does crime prevention trump privacy and liberty? When does the convenience of targeted marketing outweigh its intrusiveness? Hence we need to ask deeper questions about ethics, and specifically about the ethics of privacy, which is the subject of the next chapter. To do so, we will now take the road less travelled, right back to Ancient Greece, and then to Königsberg in 1785.

Chapter 3

Wait! Privacy? What's that?

A professional photographer and amateur voyeur, L.B. “Jeff” Jefferies is recuperating after an accident. Confined to a wheelchair in his Greenwich Village apartment, he bides his time looking out his rear window, observing his neighbours: a dancer he nicknames “Miss Torso”; a single woman, “Miss Lonelyhearts”; and a travelling salesman with a bedridden wife. One night during a storm, glass shatters and a woman screams, “Don’t!” Next the salesman can be seen acting suspiciously. Jeff promptly deduces that the salesman has murdered his wife. When a neighbour’s dog has its neck broken, Jeff infers that the salesman did that too. This is the plot of *Rear Window*, the 1954 thriller directed by Alfred Hitchcock. Eager to do something, Jeff (James Stewart) enlists the help of his girlfriend Lisa (Grace Kelly) to break into the salesman’s apartment. Inevitably, events spin out of control and into confrontation and violence, as more characters become involved, including a nurse, a detective and various neighbours. Thematically, *Rear Window* unfolds to be a film about secrets, and about privacy. It prompts several questions: What is the connection between secrecy and privacy? What sort of control ought we have over our privacy? And how does privacy link to our social relationships?

In this chapter, I address these issues by attempting to answer one simple overarching question: What is privacy? Unfortunately, the answer is less simple than the query, and has long been contested. After providing some etymological, historical and legal context, I turn to my main task, which is to define privacy. For some, this has been a matter of articulating necessary and sufficient conditions. I explore the various conditions that have been proposed in the literature and argue, ultimately, that this is an unsatisfactory approach. Privacy, with all its complexity, cannot be corralled so easily. In the literature, distinctions are drawn between personal privacy and situational privacy, between informational privacy and bodily privacy, and between the condition of privacy and the right to privacy. Some of the confusion in the literature, I suggest, stems from a conflation of these categories. Beyond this, I suggest that the best approach is to articulate a conceptual, analytic model. Hence I interrogate the relative merits of various

models, including control, access and contextual integrity, before arguing that the control model is inadequate without reference to the notion of restricted access. This leads me to support a wide conception of the access model. Once we adopt such a model, privacy begins to come into focus, just as Jeff's neighbours do through his telephoto lens.

I - A defining issue of our time – but can anyone define it?

“Privacy has become the object of considerable concern,” wrote Charles Fried in a 1968 paper now established as one of the field’s seminal essays. It was entitled, simply, “Privacy” (Fried, 1968: 475). In the half-century since, that “considerable concern” has become rather more considerable. Perhaps this would not have surprised Fried, who, on the eve of the internet’s invention, was uncannily prescient with his references to “electronic eavesdropping” and “the more insidious intrusions of increasingly sophisticated scientific devices into previously untouched areas” (Fried, 1968: 475). Accordingly, Fried has now been joined by a long list of scholars who have broached the subject, many prompted by the way in which the internet is throwing up a new set of challenges and issues. The discussion is not limited to the academy. On the contrary, the hacking of celebrity accounts, tagging on Facebook and domestic surveillance programs routinely attract mainstream headlines. In the 1960s, Fried and his peers revealed privacy as an emerging and significant subject; in the 21st century, privacy ranks as a defining issue of our time. This is in part because internet users find themselves caught in the middle of a clash between two trends: on the one hand, the internet has brought with it an ever-increasing confusion and challenge to privacy (as described in chapters one and two); on the other hand, the past century has seen a steady increase in the value placed upon privacy and the role ascribed to it, both as a general shift since 1890 and also in response to a growing awareness of the challenges posed by the internet and digital technology. The former trend has tended to work against privacy; the latter trend has tended to promote and protect privacy. Hence the concern identified by Fried now manifests in two distinct ways: there is concern about the way user privacy is being challenged on the internet; and then there is concern about privacy *per se*, which in recent decades

has become, in many jurisdictions, increasingly valued and protected by privacy regulation and protection (as discussed below).

Unfortunately, however, the topic is frustratingly intractable. “The concept of ‘privacy’ is elusive and ill-defined” (Posner, 1977: 393). “Nobody seems to have any very clear idea what it is” (Thomson, 1975: 295). “There is little agreement on the most basic questions of its scope and derivation” (Rubinfeld, 1989: 737). At times, scholarly debate has led to more confusion and obfuscation than clarity and enlightenment. As Beate Rössler writes, “The predicate ‘private’ is a complex one, which we can attribute to actions, situations, states of affairs or states of mind, places, and objects alike” (Rössler, 2004: 6). Certainly, the topic has led to a labyrinth of claims and counter-claims by philosophers, lawyers, sociologists and others. Some have given *descriptive* accounts of privacy, describing what in fact counts as private, while others have given *normative* accounts, seeking to delineate its value and the ways in which it might be protected. Some have regarded privacy as an *interest* with its own moral value, others have regarded it as a moral or legal *right* that warrants protection. Many regard privacy as somehow *valuable*, whereas some see it as a duplication of other interests and rights that is, in itself, *worthless* (see DeCew, 2015). Meanwhile, what is the distinction between *bodily* privacy and *informational* privacy? Between *personal* and *situational* privacy? (see Elgesem, 1996) What’s more, these positions and accounts themselves are at times muddled, so that, for instance, it can be unclear whether a given analysis is intended to be descriptive or normative. In this way, the task of defining privacy is sometimes enmeshed in attempts to elaborate the value of privacy. There is perhaps only one point of consensus: “One point on which there seems to be near-unanimous agreement is that privacy is a messy and complex subject” (Nissenbaum, 2010: 67; see also Tavani, 2007: 3). In short, questions are heaped upon questions. Is privacy a condition? A preference? An interest? A right? A value? Or a nuanced blend of all of these? And if it is a right, is privacy a descriptive right? A normative right? A legal right? A moral right? The aim of this chapter is to disentangle these various “privacies” in order to provide an account that aims to be ethically and legally normative.

To make sense of the confusion, one approach has been to classify the extensive scholarship into strands. Rössler identifies six types of privacy discourse, each of

which approaches the problem from a different angle, refers to a different history of privacy and emphasises a different aspect of the word's meaning. For Rössler, a first discourse considers privacy in the context of sociological and philosophical theories of the public sphere. A second discourse, also based in sociology, focuses on privacy on its own account, with particular emphasis on the "private family". A third is the discourse of feminist theory, which, as Rössler says, "has influenced the interpretations and conceptualizations of privacy in contemporary social debates more than any other." A fourth discourse, often inspired by advances in information technology, is concerned with the privacy of information. A fifth discourse is the legal discourse, which varies considerably between jurisdictions such as, say, the US and Germany. And the sixth is the philosophical discourse, in which various strands of the other discourses converge and "which since the 1960s has come to produce an independent body of philosophical literature centred upon a precise delineation of the definition and function of privacy" (Rössler, 2005: 2-4). Following this taxonomy, my aim is to draw upon several discourses, as befits the interdisciplinary nature of a thesis in philosophy and media. The focus shifts: in chapters one and two, my concern was largely privacy of information; from this point onwards, I engage primarily with philosophical and legal discourses on privacy. Here and there, however, other discourses raise their hands to contribute. This work, then, is marked by an interplay of discourses and disciplines, with a focus on the philosophical, informational and legal.

First, though, is the attempt to define privacy purely an exercise in futility? Throughout these discourses, the argument is sometimes mounted that privacy is useless, or even dangerous. In particular, feminists have long argued that privacy has been used to oppress women, given that society's public/private divide has traditionally valued and celebrated public roles, which have tended to be carried out by men. By contrast, women have been associated with (and confined to) the private and the domestic, which has tended to be undervalued. For instance, Virginia Held argues that the public/private distinction has typically privileged "the points of view of men in the public domains of state and law, and later in the marketplace, and to discount the experience of women" (Held, 2005: 87-88). On this view, the very notion of the modern public sphere depends upon the exclusion of women, with the private sphere merely the necessary and invisible foundation of the public sphere. In a media context, this has particular relevance for

conceptions of the public interest, a topic I address in chapters five and six (Lumby, 2006: 309). Further, as Anita Allen writes:

Marriage, motherhood, housekeeping, dependence, and her own moral ideals of caretaking and belonging have made many a woman's home life a privacy bane rather than boon ... Women face the problem of overcoming inequitable social and economic patterns that substitute confinement to the private sphere for meaningful privacy (Allen, 1988: 54).

Allen thus distinguishes *confinement*, or oppressive privacy, from *meaningful* privacy. On her view, privacy can be a tool of inequity. However, Allen is a privacy advocate, arguing that privacy is required by the liberal ideals of personhood, and for the participation of citizens as equals. For Allen, what is required is a liberating, rather than a restrictive, form of privacy. Hence women should be encouraged to discover individual forms of personal privacy, which will require widespread change: "Women's abilities to participate and contribute in the world as equals and on a par with their capacities are limited where laws and customs deprive them of opportunities for individual forms of personal privacy" (Allen, 1988: 53). To have privacy that is more meaningful, for instance, women should have easy access to contraception and abortion. "Decisional privacy to choose whether or not to bear a child affords fertile, younger women a valuable degree of control over the personal privacy they have at home" (Allen, 1988: 81). On this view, privacy can both dangerous and valuable. The challenge is to foster privacy as a good, and not an instrument of oppression. From the grave, the murdered wife of *Rear Window* would no doubt concur.

Others argue that privacy is derivative and insignificant. This is the view taken by privacy skeptics, otherwise known as "reductionists" (see Gavison, 1980: 422-424; and Schoeman, 1984: 209-212). Among them is James T. Moor, who argues that privacy is not a core value. After suggesting that it is conceivable that human cultures might flourish without valuing privacy at all (a minority position, as discussed in chapter four), Moor writes:

The core values are the values that all normal humans and cultures need for survival ... The core values allow us to make transcultural judgments. The core values are the values we have in common as human beings (Moor, 1997: 29).

For Moor, core values include life, happiness, freedom, knowledge, ability, resources and security. On this formulation, privacy is merely an expression of the

core value of security. Similarly, Judith Jarvis Thomson argues that a person's right to privacy is violated only if there has been a violation of another, more basic right. For Thomson, the right to privacy is in fact a cluster of rights, which always overlap with property rights or rights to bodily security. Hence, given that any privacy violation can be better understood as a violation of a more basic right, there is nothing illuminating about privacy (Thomson, 1975). Another strong reductionist argument is made by Richard Posner, who proposes that we should simply dispense with privacy altogether, given it is merely an "intermediate" value with no utility (Posner, 1977: 394).

To understand the reductionist position, we must reiterate the distinction between the condition of privacy and the right to privacy. My condition of privacy is the extent of privacy I enjoy, provided at this moment by the clothes I am wearing, the walls around me and the encryption of my emails. My right to privacy, by contrast, denotes the moral or legal entitlements to privacy I currently have. As outlined below, I have a right to privacy as prescribed, *inter alia*, under international law by the Universal Declaration of Human Rights, under Australian Law by the *Privacy Act 1988* and under the social norm that my neighbour will not walk into my home without knocking, even though my door is unlocked. Often, discussions about the condition of privacy are descriptive, whereas discussions about the right to privacy are prescriptive. An example of the former is when Parent defines privacy as "the *condition* of not having undocumented personal knowledge about one possessed by others" (Parent, 1983: 269). An example of the latter is when Reiman writes,

If ... we think that individuals ought to have others deprived of access to some of their personal affairs, whether or not a law says so, then we think that there is something like a moral right to privacy. And we will want our laws to protect this moral right by backing it up with an effective legal right (Reiman, 2004: 199).

The condition/right distinction is often subtle. In Moore's summation: "We could define privacy as *being let alone* or as a *right to be let alone*. Privacy could be cast as a *condition* that obtains or as a *right* that a condition obtains" (Moore, 2013a: 22). Often, scholars are both descriptive and prescriptive, switching between the two, sometimes with clarity, sometimes without. This can be a source of confusion (Tavani, 2007: 4). Clearly, however, Thomson's reductive dismissal of privacy is a dismissal of the *right* to privacy. She is not claiming that the

condition of privacy is a subset of the condition of security, but that the right to privacy is invariably covered by the right to security.²¹

In response, many have leapt to the defence of privacy as non-derivative and significant, arguing that privacy and security can and should be treated as distinct concerns (Inness, 1992; Bambauer, 2013). Contra Thomson, Schoeman provides the example of sound wave interceptors. Let us imagine that these come in two varieties: the first records the speech carried by the sound waves; the second converts the sound waves into usable energy but makes no record of the speech. Now suppose I have two neighbours. One trains the first device on my house, and records every one of my utterances. Another trains the second device on my house, and records every sound wave, which is then converted to energy. These two neighbours, argues Schoeman, violate my rights in profoundly different ways. The first breaches my privacy; the second does not.²² Yet for Thomson, who sees nothing distinctive about privacy, these two unique instances of recording must be categorised as morally identical. As Schoeman writes:

The suggestion here is that without reference to privacy rights specifically we shall not be able to account for the wrongfulness of certain acts consistent with the innocence of certain others. Without reference to privacy, we will not be able to draw moral distinctions which are important to describe (Schoeman, 1984: 210).

As Schoeman's reasoning shows, the right to privacy is not always covered by the right to security. For Bambauer, the distinction is yet more stark. In the context of digital technology, Bambauer writes that privacy is normative, whereas security is the mechanism by which privacy protections are put into place. The issue is:

... about clashing interests and values, and about the difficult task of choosing among them. Shifts in privacy rules nearly always burden some stakeholders while benefiting others. Rule configurations are justified by recourse to value frameworks: efficiency, distributive justice, or religious prohibitions. And these configurations describe how privacy ought to function. Security, by contrast, describes how privacy does function ... Security implements privacy's choices. Security determines who actually can access, use, and alter data (Bambauer, 2013: 676).

On the internet, Bambauer writes, security simply "mediates" privacy rights and carries them out (Bambauer, 2013: 676). On this view, privacy is invoked in a

²¹ This crucial distinction between privacy's condition and right, outlined in chapter one, comes to the fore below, and again in chapters five and six.

²² What right is violated in the second case? Perhaps none at all; perhaps my right to security; or perhaps my right to live in my home without unreasonable disturbance. The answer is unclear, but clearly my right to privacy is intact.

normative context, for its value and its right, whereas security is invoked merely as the enforcement mechanism.

In sum, there are two ways to dismiss privacy. One way is to argue that privacy does not matter, because it has no value. The strong version of this argument is that we would all be better off without privacy. A variation of this argument, outlined earlier, is the feminist position that privacy can be dangerous, and hence that it has value only in certain forms, in certain contexts. I respond to these arguments in the next chapter, where I directly address the question of why privacy matters. The second way to dismiss privacy is to argue that privacy does have value, but that privacy itself is not distinct. I respond to this argument in this chapter, above and below. In this chapter and the next, then, I depart from Moor, Thomson and Posner to sketch a view of privacy as significant and non-derivative.

II - From realm to right: etymological, historical and legal context

The idea of privacy is old. The English word stems from Ancient Rome, from the Latin verb *privare*, meaning “to deprive”, and its past participle, “*privatus*”, meaning “withdrawn from public life, peculiar to oneself, a man in private life” (TNSOED, 1993: 2359). The notion extended far beyond the Roman Empire, with an explicit distinction drawn between “private” and “public” realms for at least 2,500 years in both eastern and western philosophy. About 500 years before Christ, Confucius wrote that “a private obligation of a son to care for his father overrides the public obligation to obey the law against theft” (Moore, 2013b: 3). Contemporaneously in Ancient Greece, the distinction between public and private activity was clearly demarcated by the time of Socrates, Plato and Aristotle (Moore, 2013b: 2). In the fourth century B.C., Aristotle distinguished between the *polis*, or the public sphere of politics, and the *oikos*, or the domestic sphere of the family (DeCew, 2015: 3). In this tradition, various philosophers have continued to regard as distinct the public realm of our life as citizens and the private realm of our domestic lives within the home, where we tend to our families. This distinction, for instance, informs the philosophy of Hannah Arendt, who argued that one defining element of totalitarianism involves the destruction of private

spaces by means including the state's recruitment of family members and neighbours to spy on one another (Young-Bruehl, 2008: 52-53). In chapter one, I also noted Arendt's concern about the intrusion by the private realm into the public realm. For Arendt, private and public are determined by place: "It should be clear that the distinction between private and public depends on the locality where a person moves" (Arendt, 1977: 104). By attaching the predicate "private" to place, Arendt suggested that what is private and what is public is established, naturally and forever; she also focused on just one of the many layers of privacy's meaning (Rössler, 2004: 7). The strand of thought that distinguishes *oikos* and *polis*, private realm and public realm, comprises the first two discourses described by Rössler (Rössler, 2005: 2-3).

By contrast, the idea of *individual* privacy is not so old. In the middle ages, individual-based notions of privacy, of the sort that are now widely assumed, did not exist. The privacy of the bedroom and bathroom, with allowance for privacy in aid of sexual intimacy and bodily functions, only emerged later, alongside modern notions of the individual (Ess, 2011: 17-18). Change arrived with the spread of the Enlightenment and liberalism. In 1689, Enlightenment philosopher John Locke linked privacy with self-ownership. In the *Second treatise on government*, Locke argued that in a state of nature all the world's riches are held in common and are thus public, whereas one possesses oneself and one's own body, and can acquire property by mixing into it one's labour, thus transforming it into private property (DeCew, 2015: 3). For Locke, the public/private distinction was used to mark the limits of justified interference in personal conduct (Moore, 2013b: 4). So too for John Stuart Mill. In the 1859 essay "On liberty", Mill contrasted the realm of governmental authority with the realm of individual self-regulation (DeCew, 2015: 3). Only with the spread of the Enlightenment and liberalism, as the individual increasingly came to prominence, did privacy begin to be attached to persons, rather than just to that realm that was not the public realm.

The association of privacy and the individual culminated with the emergence of a *right* to privacy little more than 125 years ago. In 1890, an essay published in the *Harvard Law Review* was entitled, simply, "The right to privacy". Therein, US jurists Samuel Warren and Louis D. Brandeis examined US law to see whether

they could find the existence of any such right, which they described as “the right to be let alone” (Warren and Brandeis, 1890: 205). Oft-cited and highly influential, the essay is commonly regarded as the birthplace of the modern philosophy of privacy; philosopher Ferdinand Schoeman dubs it, “the first sustained and explicit discussion of privacy” (Schoeman, 1984: 202).²³ Much like this thesis, the paper was prompted by concerns about the ethics of new media. Specifically, Warren and Brandeis were dismayed by the excesses of photography and newspapers, which were becoming increasingly efficient at spreading malicious rumour. “To occupy the indolent, column upon column is filled with idle gossip, which can only be procured by intrusion upon the domestic circle,” they wrote, providing a template for analogous twenty-first century complaints against phone hacking, long-lensed paparazzi and websites devoted to gossip and scandal (Warren and Brandeis, 1890: 196).

After surveying the law of contracts, property, trusts, copyright, trade secrets and torts, Warren and Brandeis concluded that existing US law *did* offer protections for individual privacy, and particularly for the invasion of privacy engendered by the public dissemination of personal details. This general right to privacy would protect the extent to which one’s thoughts, sentiments and emotions could be shared with others. The aim was not to protect intellectual property or the items produced by individuals, but *peace of mind*. They wrote:

... a principle ... may be invoked to protect the privacy of the individual from invasion either by the too enterprising press, the photographer, or the possessor of any other modern device for recording or reproducing scenes or sounds ... [T]he principle which protects personal writings and any other productions of the intellect or of the emotions, is the right to privacy (Warren and Brandeis, 1890: 213).

Their newly-identified right to privacy, they argued, was based on a principle of “inviolate personality” which was analogous to, and connected with, the prevalent conception of one’s home as one’s castle (Warren and Brandeis, 1890: 205, 211). Before Warren and Brandeis, the *condition* of privacy had been at issue, whether for Confucius, Aristotle or Mill. Indeed, it stayed at issue for Arendt. After 1890, while the condition of privacy remained relevant, it became complemented by the

²³ Admittedly, Warren and Brandeis did not invent the right to privacy, nor even the phrase the “right to be let alone”. Rather, their analysis of the common law in search of a right to privacy gave the principle shape and form (Gavison, 1980: 423-424). As such, the modern right to privacy is conventionally traced back to their essay, and Warren and Brandeis are commonly regarded as its inventors.

notion of privacy as an individual's moral or legal entitlement. Hence any comprehensive account of privacy must look further than just the notion of a private "realm", and must address both privacy's condition and right.

The paper's impact on the law is hard to overstate. As early as 1954 it was described as "perhaps the most famous and certainly the most influential law review article ever written" (Nimmer, 1954: 203). It did, however, take several decades for privacy to gain a solid legal footing in the US, as elsewhere. In 1928, now as a Supreme Court judge, Brandeis handed down a judgment in which he reiterated the phrase "the right to be let alone", describing it as "the most comprehensive of rights and the right most valued by civilized man" (*Olmstead v United States*, 277 U.S. 438, 478 (1928), quoted in Christians et al., 2012: 95). It was, however, a dissenting judgment. Then, in the 1965 case of *Griswold v Connecticut* (381 U.S. 479), a majority of the Supreme Court came to agree with Brandeis, emphatically announcing the right's arrival with a 7-2 verdict. In *Griswold*, it was found that the Constitution protected a right to privacy, even though the word does not appear in the Bill of Rights (DeCew, 2015: 6). Rather, wrote Justice William O. Douglas, the right was to be found in the "penumbras" and "emanations" of various constitutional provisions. These included the Fourth Amendment to the Bill of Rights, which defends "the right of the people to be secure in their persons, houses, papers, and effects, against unreasonable searches and seizures" (LII, 2016). Subsequently, US courts have solidified the legal right to privacy, invoking it to overturn a ban on interracial marriage, to allow people to possess obscene materials in their own homes and to dispense contraception (DeCew, 2015: 6-9). In 1973, in the enduringly controversial case of *Roe v Wade*, the US Supreme Court upheld a woman's right to have an abortion, finding an implied right to privacy in the Fourteenth Amendment, which provides that, "No State shall make or enforce any law which shall ... deprive any person of life, liberty, or property, without due process of law ..." (LII, 2016).

To supplement the case law, a series of statutes have been passed to protect privacy in the US. At a federal level, statutes cover everything from the privacy of medical records to the privacy of students (CDT, 2008: see chapter six). At a state level, California has passed a law against revenge porn (see chapters five and six). Whereas *Roe v Wade* is an instance of *bodily* privacy, revenge porn laws are an

instance of *informational* privacy, and both are far removed from the *oikos/polis* distinction drawn by Aristotle and Arendt. Further, both *Roe v Wade* and revenge porn laws are typical of the manner in which privacy has been finding its way into law in the US and elsewhere, by courts and parliaments willing to find privacy protections in instruments such as the constitution in response to specific issues.²⁴ Since 1890, the legal protection of the right to privacy has tended to be inventive and interstitial (Gavison, 1980).

In Australia, the right to privacy is not as clearly articulated, regulated or protected as in the US. The Australian Constitution has no Bill of Rights to provide for privacy, either explicitly or in its penumbras. Rather, Australia protects privacy in an even more piecemeal way, via a complex of federal and state legislation, local regulation and case law. Tellingly, of the many laws on the topic, not one contains a definition of privacy. As a result, the Australian Law Reform Commission (ALRC) has resorted to conducting inquiries into privacy by invoking the contextual use of the term rather than any legal definition (ALRC, 2007). In Australia, no tort of invasion of privacy exists, but the High Court has left open the door for the development of such a cause of action in common law (ALRC, 2007: §5.12). In its 2014 report *Serious invasions of privacy in the digital era*, the ALRC recommended that a new tort be created to provide a statutory cause of action for invasion of privacy (ALRC, 2014: chapter 4). I return to this issue in chapter six, where I argue for the enactment of such a cause of action. In the absence of such a tort, privacy is protected by a messy assortment of federal and state laws. Federal laws include the *Telecommunications Act 1977*, *National Health Act 1953* and also the *Privacy Act 1988*, which establishes principles that constrain the handling of personal information by government agencies and large businesses. It does not, however, provide protection against the actions of small businesses and other individuals. Recently, local regulations have also been passed to help protect privacy. In 2015, one Sydney council, spurred on by government inaction at state and federal levels, was the first to ban drones in parks and public spaces (Gair, 2015).

²⁴ In the US, the “constitutional” right to privacy is regarded as distinct from “informational” privacy (Bloustein, 1964: 962; DeCew, 2015).

In Australia, miscellaneous remedies also exist at common law. Equitable actions for breach of confidence have been recognised by Victoria’s Supreme Court of Appeal in *Giller v Procopets* and the Western Australian Supreme Court in *Wilson v Ferguson*. Both were cases of revenge porn, in which monetary compensation was awarded for misuse of personal information (ALRC, 2014: §12.14-12.22; Chighine, 2015). It has been argued that several Australian courts, keen to protect privacy, are actively compensating for a lack of a tort of privacy by extending remedies for breach of confidence (Gatford, 2015). This approach would follow recent developments in UK case law (Mo, 2017: 87-89; see chapter six). Another common law remedy is trespass, which gives the occupier of a property a significant degree of control over who comes onto her property. In some jurisdictions, such common law is supplemented with legislation, such as section 4 of the *Inclosed Lands Protection Act 1901* in the Australian state of New South Wales, which provides, “Any person who, without lawful excuse (proof of which lies on the person), enters into inclosed lands without the consent of the owner, occupier or person apparently in charge of those lands, or who remains on those lands after being requested by the owner, occupier or person apparently in charge of those lands to leave those lands, is liable to a penalty ...” (ALII, 2016b).

In European law, by contrast, privacy has been far more clearly articulated and protected than in Australia or the US.²⁵ Rather than interstitial, Continental law has sought to be general and broad, with the European Court of Human Rights providing explicit protections based on Article 8 of the European Convention on Human Rights, which states: “Everyone has the right to respect for his private and family life, his home and his correspondence.” The limits of Article 8 have been regularly tested. In a 1992 case, the court construed the phrase “private life” to encompass more than simply a private realm. In *Niemietz v Germany*, authorities searched the premises of a lawyer to find the identity of a man who had written insulting letters anonymously, an offence in Germany. The lawyer complained that the search interfered with his private life, and the court agreed. Interpreting the phrase “private life”, the court found:

²⁵ On one account, the differences in privacy law in the US and Europe arise because two different legal systems with two different histories are protecting different values. In the US, privacy primarily protects a liberty interest, whereas on the Continent, privacy laws are based on French and German notions of personal honour (Whitman, 2004).

... it would be too restrictive to limit the notion to an ‘inner circle’ in which the individual may live his own personal life as he chooses and to exclude therefrom entirely the outside world not encompassed within that circle. Respect for private life must also comprise to a certain degree the right to establish and develop relationships with other human beings (*Niemietz*, 1992).

The court thus held that Article 8 protects the right to make and maintain relationships with other people, including at work, given that it is at work that people often have “a significant, if not the greatest, opportunity of developing relationships with the outside world” (*Niemietz*, 1992). Hence the court’s conception extended beyond the *oikos* of Aristotle to recognise that privacy can prevail outside the private realm. European courts recognise that European privacy protections are more stringent than US protections (*Schrems*, 2015). European privacy protections tend to fall under the umbrella of “data protection”, and hence within the parameters of what is known as “informational privacy”. This notion, for instance, underpins the newly recognised “right to be forgotten”, which, a court held in 2014, entitled a Spanish man to have certain adverse hyperlinks concerning his past financial status to be removed from Google’s returned search results (Kranenborg, 2015: see also chapter six). By November 2015, Google had received nearly 350,000 right to be forgotten requests for 1.2 million links to webpages. Of those, 42 per cent of URLs were removed from search results and 58 per cent were retained (Calpito, 2015). When it comes into effect in May 2018, the General Data Protection Regulation, or GDPR, will go even further to articulate and protect privacy in Europe, including the right to be forgotten.²⁶

Globally, a vast labyrinth of privacy laws exist, with most countries protecting privacy in their constitutions (Greenleaf, 2015; Solove, 2008: 3). Brazil’s Constitution provides that “the privacy, private life, honor and image of the people are inviolable”; South Africa’s constitution prescribes that “everyone has the right to privacy”; and South Korea declares that “the privacy of no citizen shall be infringed” (Solove, 2008: 3). The Qatari Constitution provides that the sanctity of human privacy is “inviolable” and that no interference into the privacy of a person is allowed unless permitted under the law. Among other provisions, the Qatari Penal Code also prohibits the spreading of news, photographs or

²⁶ National and international privacy law, including Europe’s right to be forgotten and the GDPR, is further discussed in chapter six.

comments related to secrets of the private life of families or individuals, and prohibits intrusion into a person's private life without their consent and not in accordance with the law (Abokhodair et al., 2016: 68). In international law, meanwhile, the concept of a right of privacy became increasingly established during the twentieth century, most notably in Article 12 of the Universal Declaration of Human Rights, adopted in 1948: “No one shall be subjected to arbitrary interference with his privacy, family, home or correspondence, nor to attacks upon his honour and reputation. Everyone has the right to the protection of the law against such interference or attacks” (UN, 1948). This forms the basis of Article 8 of the European Convention on Human Rights, discussed above. In this formulation, privacy is not defined, although an implied meaning is evoked via references to family, home, correspondence, honour and reputation. The same wording appears in Article 17 of the International Covenant on Civil and Political Rights, which came into force in 1976 (and which inspired, *inter alia*, the wording of Australia’s *Privacy Act 1988*).

From the foregoing, we can draw several conclusions. First, with the spread of the Enlightenment and liberalism, there has been an increasing association of privacy with the individual. Second, in recent centuries, and especially in recent decades, privacy has increasingly come to be seen as valuable and thus worth protecting. Third (and tied in with these first and second points), the right to privacy, since emerging in 1890, has become established as a significant ethical and legal principle, explicitly recognised by the law of numerous jurisdictions, as well as internationally. Fourth, the definition of privacy is far from settled. Indeed, the legal definition is often implied and contextual. And fifth, legal protections take many forms. In Australia, privacy has limited protections from a patchwork of laws, and “privacy” itself is never defined. In the US, privacy is emerging interstitially from landmark judgments relying on implied provisions, as well as specific legislative instruments. In Europe, privacy tends to be more clearly defined and better protected, codified in the GDPR and elaborated in a weighty body of case law that links privacy to, among other things, family life and the right to establish relationships. And in Brazil, South Africa, South Korea and Qatar, privacy enjoys explicit constitutional protection. Legal protections thus range from the implicit and interstitial to the explicit and expansive. Privacy protections also target different *types* of privacy. One approach involves the

protection of the private realm, as seen in the law of trespass; another involves the protection of privacy that attends the individual whether or not she is in a private space, as recognised in *Niemitz v. Germany*; another involves informational privacy, as in Europe's newly-recognised "right to be forgotten"; and yet another involves bodily privacy, as in the landmark abortion case of *Roe v Wade*.

Clearly, privacy is contested. It is also expanding. It has extended from *oikos* to the individual to a right to the body. What is understood by "private" in the 21st century is much more broad and encompassing than in the time of Confucius and Aristotle; these days, "private" and "privacy" attach themselves with more subtlety and elasticity to a remarkably wide range of exchanges, activities and attributes. At the same time, as I have described in chapters one and two, privacy is being challenged and confused by our internet interactions. Privacy, in this sense, is under threat. Here, then, are two conflicting trends: just as it is being more recognised and protected, privacy is increasingly being challenged (and confused) by our use of technology. Just as it is flourishing, then, it is also being hamstrung. More specifically: just as the *right* to privacy is increasingly being recognised and protected, the *condition* of privacy is increasingly coming under threat. Internet users, it would seem, have less privacy even as they acquire more right to it.

III - Conceptual accounts

Privacy is the semantic equivalent of a double-jointed yogi. Its flexibility is remarkable. My diary? My sexual preference? The god I worship? My email correspondence? My medical records? The conversation I have with a friend, even when we're seated in a public café? My body? These aspects of my life, and many more, are covered by the notion of privacy (Rössler, 2004: 6). In Ancient Greece, privacy was confined to the realm of the *oikos*. As such, the term was limited and precise. Today, that meaning persists, but privacy is also associated with individuals. It extends into their relationships and follows them into public spaces. Often, it is protected in law in the form of the right to privacy. Today, then, the term "privacy" has transcended its neat and quarantined origins to grow complex and unwieldy.

What I hope to do now is to make some sense of the perplexing flexibility that has come to attend the notion of privacy since the time of John Stuart Mill.

Specifically, what I am seeking is not a dictionary definition, but a *conceptual account* narrow enough to be meaningful, yet broad enough to cover both the scope and subtlety of modern interpretations of privacy. I do so in the post-Mill tradition of liberalism, which itself accommodates pre-existing accounts of private and public realms. As Rössler notes, the distinction between private and public is both fundamental and constitutive for the political and philosophical school of liberalism, given that the private/public distinction seeks to protect individual freedom and autonomy in the face of impermissible interference from the state (Rössler, 2005: 10). Further, it is important to note that we are discussing *individual* privacy. That is, privacy as it pertains to individual human beings. Privacy can also pertain to couples, or families, or groups; but the focus of this thesis is individual privacy. Despite this focus, however, I will not be presenting merely a conventionally liberal account. Instead, I seek to expand upon the post-Mill tradition of liberalism by arguing that individual privacy matters for the individual in question, but also for that individual's relationships, given that we are all beings-in-relation, constituted by the various links that embed us in society. In this way, my approach is not merely individualistic, as I build in chapter four towards a notion of relational privacy. The aim is to transcend an atomistic account and to develop a conception of privacy as relational, in a manner analogous to relational conceptions of autonomy (Mackenzie and Stoljar, 2000: 3-31). The notion of relational privacy, I argue, is bound up in notions of relational autonomy.

Of course, many others have attempted to articulate a conceptual definition of privacy in a liberalist framework. Their essays have yielded remarkably divergent results. Adam Moore argues that conceptions of privacy fall into six categories: the right to be let alone; secrecy; intimacy; control over information; restricted access; and privacy as a cluster concept (Moore, 2013b: 6). Conversely, Herman Tavani organises the classic philosophical and legal theories of privacy into four categories: the nonintrusion; seclusion; limitation; and control theories of privacy. However, he also notes that several alternative schemes for categorisation exist (Tavani, 2007: 4). A more straightforward approach is to identify three pre-

eminent conceptual models: the first defines privacy in terms of control (Fried, 1968; Elgesem, 1996); the second defines privacy in terms of restricting access (Gavison, 1980; Tavani, 2007); and a third, more recent strand is a form of definitional agnosticism that attends the theory of “contextual integrity” (Nissenbaum, 2010). In what follows, I explore various conceptual accounts, with a particular focus on control, access and definitional agnosticism.

According to Rössler, there are two predominant conceptions in common usage. In one, privacy resembles an onion. On this view, there exist layers of privacy, with whatever is most intimate and personal comprising the innermost layer (Rössler, 2004: 6-7). Here lies bodily privacy. In the second layer can be found family and other intimate relationships. Moving outwards, another layer represents community. On the very outside layer is the state, the skin of the onion that is fully public. One version of this conception is the traditional *oikos/polis* distinction, in which privacy pertains to *place*; here, privacy prevails in a “private realm”, “private space” or “private domain”. Indeed, 2,400 years after Aristotle, privacy and place remain closely related. Along these lines, arguments have been advanced to say that private spaces need greater protection. Iris Marion Young, for instance, proposes that privacy theories have paid insufficient attention to supporting the condition of privacy by guaranteeing personal space. This personal space, she argues, involves the value of home (Young, 2004: 168). The commonly-drawn link between privacy and place is also evident in the way homes are considered private, and more so in the way the bedroom and the bathroom are considered private spaces *par excellence* (Reiman, 2004: 198; Ess, 2011: 17-18). Privacy thus extends to the activities that commonly occur therein. Sleeping is considered private. Dressing, undressing and engaging in physical intimacy are considered private. Going to the toilet is considered private. Even though privacy’s purview has expanded markedly, it is true that there remains a significant link between privacy and place. Reading the newspaper, I would be surprised and affronted if a man peered in at my family through our living room window.

As we have seen, however, privacy is not just about place. Sometimes it attaches to the individual. Sometimes privacy follows the individual into public spaces, as with certain phone conversations held in public. And sometimes it seems to attach

to the *situation*. If a child relieves herself in public, one would expect people not to stop and stare. The same is true of pregnant women, who are in many jurisdictions exempt from prohibitions on public urination. This brings us to Rössler's second common usage conception, in which privacy attaches itself to certain actions and decisions (Rössler, 2004: 6-7). On this view, dubbed the "dimensions" model, "private" refers to particular actions or decisions, no matter where they physically occur. Hence going to church and what I choose to wear are both private affairs. "Here the concept of the private describes a protected sphere of individual action and responsibility, where individuals are able to act independently of the decisions or interferences of other people, of public authorities and institutions" (Rössler, 2004: 7). Rössler's first common conception of privacy invokes the physical, including spaces, realms and even the body as the onion's innermost layer; her second conception, the dimensions model, involves actions and decisions. This division seems to be supported by Dag Elgesem, who proposes that we must be clear to distinguish between *situational* privacy and *personal* privacy (Elgesem, 1996: 48). As Rössler notes, the *spaces* account and the *dimensions* account (which also often links privacy to place) sometimes operate at cross-purposes.

However, as Rössler and Elgesem show, neither the onion model nor the dimensions model on its own tells the whole story. Instead, let us surmise a connection, as several philosophers have done, between privacy and intimacy, broadly-conceived, so as to include notions of seclusion, vulnerability, the body and the personal. Tellingly, a common euphemism for genitalia is "privates". Several scholars argue that privacy simply could not exist without intimacy (Gerety, 1977; Gerstein, 1978; Inness, 1992). I address many of these arguments in chapter four, which concerns the value of privacy; here, however, I am concerned with the definition of privacy. In this regard, Fried argues that privacy is necessary for one's development as a moral and social individual who is able to trust, love and befriend. These social bonds can only be built by a judicious and selective sharing of intimacies (Fried, 1968: 484). Gerstein also argues that privacy is necessary for intimacy, and intimacy is necessary in our communications and relationships for us to fully experience our lives. Such intimacy can only occur without intrusion or surveillance, so that we can act spontaneously and shamelessly (Gerstein, 1978: 77). And for Inness, intimacy is

the defining feature of privacy invasions. Intimacy, for Inness, is based on intention, not behaviour, and draws its meaning from love, liking or care. Privacy is what enables a person to be intimate with behaviour or information in a way that fulfills her need to love and care (Inness, 1992).

Even if we accept a link, however, the problem is that not all forms of privacy involve intimacy, and not all forms of intimacy involve privacy. The privacy that I demand for my medical records, for example, may or may not be related to intimacy. Details of gynaecological procedures are intimate; but details of a tetanus shot are not, and yet I may nonetheless want these details to remain private. The privacy of medical records is not always a function of the intimacy of the details they contain. Rather, they may be a function of the potential for embarrassment, my future job prospects or some ill-defined unease with having the world know my medical ailments and history. Conversely, I can have an intimate moment in full public. My wife and I can hold hands or kiss in way that is intimate, even though we do not require or expect any privacy for our exchange of affection. Privacy and intimacy are sometimes closely connected. Their overlap is significant. However, intimacy on its own is an insufficient notion by which to define privacy.

The same is true of secrecy. The plot of *Rear Window* suggests privacy and secrecy might be one and the same, but can my bedroom and bathroom activities be described as secret? Not typically. As a human being, I am, to a certain extent, constrained by biology. There is little *secrecy* about my activities in those rooms. And yet I still hope that privacy prevails, so that, unobserved, my dignity will be respected. Sometimes, then, there can be an expectation of privacy that has nothing to do with secrecy. And sometimes there can secrecy, but no expectation of privacy. Think of a gambling addict. She plays the pokies and is threatening to bankrupt her family. She keeps her addiction secret, because she knows her husband would be deeply concerned, her employer might entertain thoughts of termination, and her friends would intervene in some way she would rather avoid. Arguably, she should have no expectation of privacy about this information, particularly given the direct harm she is doing to herself, her husband, her children and others. Her addiction, it would seem, is secret, but not private.

What, then, about relationships? Alongside those who define privacy in terms of intimacy are those who define privacy in terms of social relations. Sometimes it is difficult to prise apart the two groups. For Fried, privacy is fundamental for an individual to develop into a moral and social being who can forge intimate relationships that involve respect, love, friendship and trust. Without privacy, there is no intimacy; without intimacy, there is no respect, love, friendship and trust. Building on the Kantian notion of respect for persons, Fried thus proposes that privacy is a key to our integrity as persons, and that a threat to privacy is a threat to our integrity as persons. For Fried, then, privacy is integral for both intimacy and relationships (Fried, 1968). Along similar lines, Schoeman argues that privacy provides a way to control intimate information about oneself, which in turn has benefits for the development of one's own personality, and also for the cultivation of friendships with others (Schoeman, 1984). Even more forcefully, Rachels claims that healthy social relations depend on privacy. He argues that privacy is necessary if we are to maintain relationships, and that this is true for both non-intimate relationships and intimate relationships. Privacy is how we determine who knows what about us and who has access to us so that we can make and maintain all our social relations (Rachels, 1975).

Again, however, privacy cannot be defined merely by reference to social relations. Many social relations have nothing whatever to do with privacy. A doctor or bureaucrat may or may not know all manner of private, personal details about me, but our relationship might be distant and impersonal either way. Simply, the giving or withholding of privacies is in itself not enough for the formation of social relations. Apart from privacy, other values, including affection and respect, play a part in the formation of such ties. Jeffrey Reiman, for instance, argues that relationships are much more a function of how much people care for one another than how much they know about one another (Reiman, 2004: 198; see chapter four). Further, as James T. Moor points out, there are people who do not want relationships. Is privacy therefore irrelevant to them? And what about people who do not need privacy to form a variety of relationships? Privacy must be defined and justified with regard to more than just the relationships it enables (Moor, 1997: 28). Like intimacy, social relations may provide one ingredient in an adequate definition of privacy, but they cannot provide a complete and satisfactory definition.

In the face of these difficulties, some have argued that attempts to define privacy have tended to be more of a hindrance than a help. As Daniel J. Solove writes:

The quest for a traditional definition of *privacy* has led to a rather fruitless and unresolved debate. In the meantime, there are real problems that must be addressed, but they are either conflated or ignored because they do not fit into various prefabricated conceptions of privacy ... In this way, conceptions of privacy can prevent the examination of problems (Solove, 2007: 759).

One response is to abandon the hunt altogether. This is the approach of Helen Nissenbaum, who describes privacy as a “conceptual morass”, and then dodges the morass with a deft sidestep.

Attempts to define it [privacy] have been notoriously controversial and have been accused of vagueness and internal inconsistency – of being overly inclusive, excessively narrow, or insufficiently distinct from other value concepts. Believing that one must define or provide an account of privacy before one can systematically address critical challenges can thwart further progress (Nissenbaum, 2010: 2).

Without a definition, Nissenbaum proposes “contextual integrity” as the key to unlock the issue of online privacy. Contextual integrity, founded on definitional agnosticism, is based on two principles: first, people engage in activities in a “plurality of realms”; second, each of these realms has a distinct set of norms governing it. In other words, privacy hinges upon context. It doesn’t hinge on the information in question, but on the context in which that information is shared.²⁷ Online engagements usually have offline analogues, she writes, and the appropriate norms for online engagements can usually be determined by identifying the correct offline analogue (Nissenbaum, 2011: 38-40). The contextual integrity approach, then, has its focus on the *granularity* of privacy, and Nissenbaum’s project is to establish which privacy norms should apply in a given online milieu. The approach has proven highly influential (Tene and Polonetsky, 2015: 88; Shvartzshnaider et al., 2016). Researchers are applying it widely, including to Facebook (Hull et al., 2011).

However, as Nissenbaum admits, not every online interaction has an offline counterpart. One example is a search engine such as Google. Here, Nissenbaum has proposed that the most appropriate analogue is of a library. Users of Google should thus be able to expect that they are able to search for information privately

²⁷ In response, I agree that privacy norms can depend upon context. However, this does not mean that what privacy *is* depends on context. My argument below is constructed on this premise.

and without records being kept, just as they are able to do in a physical library (Nissenbaum, 2011: 40-41). This analogue is problematic. First, libraries tend to be public institutions, whereas Google is a private company. Second, Google is company whose very revenue model is, as we have seen, largely predicated on the collection of personal data. Third, the purpose of Google's search engine is to help users locate information in the vast repository of data that is the web; libraries, by contrast, *are* the repositories of data. Fourth, Google is vast. Its scale dwarfs any pre-existing library. Fifth, Google is a digital shapeshifter that personalises itself for every user, offering diverging search results and highly tailored advertising based on each individual's detailed profile (Pariser, 2011: 1-3; see chapter two). And sixth, Google's parent company, Alphabet, doesn't just offer a search engine, but also offers email, maps, music and much more, then links all the user data collected by those services. This is not to suggest that the theory of contextual integrity is doomed, merely that it faces challenges. In any case, my larger point is that Nissenbaum's approach patently doesn't assist us in our quest to find the best possible conception of privacy. My approach, by contrast, is to formulate a definition before applying a normative principle. If it's a definition we want, contextual integrity cannot help.

IV - The control and access models

The two prevailing accounts of privacy are the control model and the restricting access model (Elgesem, 1996: 48-51; Reiman, 2004: 197-198). The control model has been articulated by a comprehensive body of scholarship spanning 50 years. In 1967, Alan Westin published *Privacy and Freedom*, which arguably remains the "inaugural treatise of the present generation of privacy scholarship" (Allen, 1988: 42). Therein, Westin argued that privacy should be defined as control over information about oneself. Hence privacy 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" (Westin, 1967: 7). Further, Westin identified people's right "to control, edit, manage, and delete information about them." On this view, the bounds of privacy are determined by each individual (or group). This view was also favoured by Charles Fried, who wrote in 1968,

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. To refer for instance to the privacy of a lonely man on a desert island would be to engage in irony. The person who enjoys privacy is able to grant or deny access to others (Fried, 1968: 482).

Note that Fried incorporates the word “access”. Prising apart the control and access models is not always easy, as we shall see. However, at the heart of Fried’s definition is the notion of control, and the way it enables us to determine our privacy by granting or denying access. Fried proceeded to further refine his definition by stating that privacy is not just about controlling the quantity of information, but also about “modulations in the quality of the knowledge.” Hence you may willingly allow an acquaintance to know you are sick, but may feel your privacy has been compromised if that acquaintance learns of certain specific symptoms of your sickness (Fried, 1968: 483).

Variations followed. In 1983, William Parent proposed a narrower definition of control over information. Whereas Westin and Fried were addressing the right to privacy, Parent addressed the condition of privacy, which he defined as a moral value for people who prize individuality and freedom. For Parent, the condition of privacy equated to “undocumented personal information”. That is, privacy exists when personal information (being the facts most people choose not to reveal) is not on the public record. Once information is on the public record, he argued, there can be no invasion of privacy (Parent, 1983). This is a particularly problematic proposition in the context of the internet. In Panopticon 2.0, all data is, at least potentially, public. This suggests that, on Parent’s account, there can be no privacy, and hence no invasion of privacy, on the net. To be fair, Panopticon 2.0 is a theoretical construct; as I showed in chapter two, limits to watching do exist. Not all data is public. For users, there is a distinction between private family photos stored in the cloud and comments made in public *fora* such as Twitter. Still, the private domain is vulnerable to access by hackers, companies and governments. Parent’s account appears far too narrow, particularly for the internet. Dag Elgesem takes another approach, but also identifies condition and control. Elgesem discerns two distinct senses of privacy: *situational* privacy, which is the state in which a person can find herself, such as inside her own home; and *personal* privacy, which “is more like a property, namely, the property of having control over the flow of personal information ... To have personal privacy

is, on my account, to have the *ability to consent* to the dissemination of personal information” (Elgesem, 1996: 48). For many scholars, then, individual privacy is indelibly linked with individual control, and sometimes this control is expressed in terms of the ability to consent. The connection between control and consent forms the subject of chapter five; for now, note merely that control forms the *sine qua non* of many modern definitions of privacy.

Control, it seems, reveals something fundamental about privacy, but is it really necessary? If I am a man who wants to go swimming, I can choose how much of my body to reveal by bathing in a full-body wetsuit, in board shorts or in the sort of truncated briefs known in Australia as “budgie smugglers”. (“Budgie” is short for “budgerigar”, a small bird.) If I am a woman, I can wear a wetsuit, one-piece, bikini or g-string. In this choice, I have control. However, the notion of control also *lacks* something. Some bathing suit options are closed off to me by regulation and by convention. Witness, for instance, the slightly different alternatives considered acceptable for men and women. More tellingly, consider the law. In Sydney, laws and regulations prohibit me from swimming or sunbathing naked, except at a handful of designated beaches (e.g. ALII, 2016a). Conversely, on certain beaches in France, the “burqini” swimsuit was banned in 2016, leading to curious scenes of fully-uniformed police ordering Muslim women to remove clothing while by the sea (Dearden, 2016). In other words, it seems that both my condition of privacy and my right to privacy are not entirely up to me. To some extent, they are out of my control, given they are subject to laws, regulations and social norms.

Unsurprisingly, then, numerous arguments have been mounted against the control model. Addressing the *condition* of privacy, Moor cites the example of a priest receiving confession. The confessor, Moor argues, has no control whatsoever over what the priest will do with the information once the confession is finished (Moor, 1990: 78).²⁸ And for Judith Jarvis Thomson, control and the *right* to privacy clearly come apart:

²⁸ In his counter-argument, Elgesem notes that the confessor has a certain control by virtue of a well-established set of norms that bind the priest to confidentiality (Elgesem, 1996: 48-49; Glancy, 1979: 2-3). I suggest this is stretching the limits of “control”.

If my neighbor invents an X-ray device which enables him to look through walls, then I should imagine I thereby lose control over who can look at me: going home and closing the doors no longer suffices to prevent others from doing so. But my right to privacy is not violated until my neighbor actually does train the device on the wall of my house (Thomson, 1975: 304).

As Thomson shows, my right to privacy is only violated if I am being spied upon. Until such spying occurs, my neighbour merely has the *potential* to violate my right to privacy. Thomson shows that the right to privacy can prevail, even when control has been lost. However, Thomson does not address the condition of privacy. I would suggest that my condition of privacy *is* altered as soon as my neighbour invents the device, whether or not it is ever used. This is more obviously so if I know about the invention of the device. It is not that my condition of privacy has *actually* been compromised by the device's invention, but merely that it has *potentially* been compromised. Suddenly, my walls no longer shield me against my neighbour's eyes, and I am potentially as exposed inside my home as outside it. And if many people owned such devices? Clearly, a world in which such x-ray devices exist is different from a world without such devices. This ray gun makes me vulnerable to arbitrariness. In sum, the technology itself does not impact my right to privacy; my right endures, and is only violated once the device is trained upon me. As such, I lose control but retain my right. By contrast, my condition of privacy is instantly *altered*, albeit not necessarily *compromised*, as soon as my neighbour invents the device.

By analogy, Thomson's argument extends neatly to the internet, where user privacy is often potentially, if not actually, compromised. Accordingly, the internet *per se* also constitutes a persuasive counter-argument to the control model. Online, as we have seen, users appear to have control over their privacy, including in the way they can adjust their privacy settings. This control, however, often turns out to be deficient or illusory. Online, control is often incomplete and compromised. In this regard, Nissenbaum outlines the "transparency paradox", or "transparency dilemma", to describe the way notice-and-consent models of privacy seem doomed to fail on the net (Nissenbaum, 2011: 36). If the privacy settings of a website are made simple, brief and intelligible for the average user, those settings cannot hope to capture the complexity of online information flows. If, however, the privacy settings are sufficiently nuanced to capture the complexity of online flows (if this is even possible), then they will be so lengthy

and labyrinthine that very few, if any, users will ever read them, let alone understand them (Nissenbaum, 2011: 36). One review of the Amazon Kindle terms and conditions (which, admittedly, are not limited to privacy) found they would take the average user nine hours to read (Cormack, 2017). In this context, privacy settings are necessarily inadequate, and user control is often illusory. In chapter two, I described the data mining giant Acxiom. In its privacy principles, Acxiom acknowledges the importance of control: “Acxiom recognizes that individuals should be informed as to how information about them is used and should have choices about the dissemination of that information” (Acxiom, 2016). Yet it seems few people know Acxiom exists, let alone that it deals in their data. Online, user control is often touted as a core principle, encapsulated in terms such as “informed”, “choice” and “consent”; the reality of this control, however, is often flawed and usually complicated. In one study, researchers showed that, paradoxically, more control can lead to less privacy, because people take more risks with their data if they feel more protected (Brandimarte et al., 2013). Apparent control does not mean actual control. What’s more, as Thomson’s x-ray device shows, I can lose control but keep my right to privacy. And, as my swimsuit example shows, both the condition of privacy and the right to privacy are about more than just control. Sometimes, both condition and right are about externally-imposed limitations, such as laws against public nudity, or against the modesty of burqinis.

To put it another way: my privacy is not just about what *I* want. Indeed, my privacy is not just about *me*. Laws prescribe nudity at Bondi Beach and laws ban burqinis on the Riviera. Are these laws enacted to protect the modesty and dignity of the bathers in question? Perhaps. Clearly, however, these laws are also designed largely with respect to other bathers. In the Australian state of New South Wales, nude bathers can be sentenced to jail under section 61N of the Crimes Act for committing an “act of indecency”, and the penalty is greater if that act is “towards a person under the age of 16 years” (ALII, 2016a). Individual privacy is not just about the individual whose privacy is at issue, but about others too. The law sometimes asks: if I renounce my privacy, will that adversely affect others? If so, which others? And how badly? We do not simply enable an individual to have all the say about the limits of her privacy, because the rest of society will want a say as well. As noted above, I am building towards an account

of privacy as relational, founded on a notion of autonomy as relational. “Its [relational autonomy’s] starting point is the individual as situated in, shaped, and constrained by her socio-relational context in all its complexity; that is, its starting point is non-ideal agents in a non-ideal world, characterized by social oppression, injustice, and inequality” (Mackenzie, 2014: 23; see chapter four). Similarly, relational privacy begins by acknowledging that an individual exists in a socio-relational context. This involves the recognition that my privacy matters not just for me, but also for others. It also involves the recognition that privacy isn’t just about control. In part, my privacy is, and ought to be, set by others.

Control, it seems, is a core component of privacy. However, privacy is demonstrably sometimes beyond my control. The limits of my privacy, it seems, are sometimes set by someone other than me, and sometimes *should* be set by someone other than me. This brings us to the restricted access model, which holds that privacy can be defined as restrictions upon access to oneself and to information about oneself. These restrictions might involve control. That is, they might be imposed by the person herself. Alternatively, they might be externally-imposed. *Prima facie*, such an approach seems to dovetail with common usage, and with dictionary definitions of privacy, which rely heavily on notions of withdrawal, non-intrusion, seclusion and secrecy. *The New Shorter Oxford*, for instance, defines privacy as: “1. The state or condition of being withdrawn from the society of others or from public attention; freedom from disturbance or intrusion; seclusion ... 2. Absence or avoidance of publicity or display; secrecy ...” (TNSOED, 1993: 2359). Nowhere does this definition include any notion of control. For scholars including Sissela Bok and Ruth Gavison, access is key. In 1982, Bok set a template with her definition: “Privacy is the condition of being protected from unwarranted access by others – either physical access, personal information, or attention” (Bok, 1982: 10-11). Along similar lines, Gavison wrote that the extent to which others have access to information and physical access to us will determine how much privacy we have. As such, we have perfect privacy only when we are perfectly inaccessible to others. For Gavison, Fried is wrong when he says that it would be ironic to talk of a man on a desert island as having privacy. His inaccessibility gives him tremendous privacy, via his solitude. Indeed, Gavison argued that privacy can be gained in three ways: secrecy (when

others don't have information about x), anonymity (when others don't pay attention to x) and solitude (when others don't have physical access to x).

These three elements of secrecy, anonymity, and solitude are distinct and independent, but interrelated, and the complex concept of privacy is richer than any definition centered around only one of them (Gavison, 1980: 428-429).

Advocating an explicit legal commitment to privacy, Gavison's paper has proved extremely influential. In 1990 James H. Moor wrote:

The core idea of restricted access accounts is that privacy is a matter of the restricted access to persons or information about persons ... By my definition, an individual or group has privacy in a situation if and only if in that situation the individual or group or information related to the individual or group is protected from intrusion, observation, and surveillance by others (Moor, 1990: 76, 79).

And in 1988, Anita Allen elaborated a restricted access view broader than Gavison's to allow for the feminist recognition that women experience losses of privacy unique to their gender (Allen, 1988).

The access model holds that privacy concerns the access of others to one's body and to data about themselves, and sometimes to one's space or possessions. Sometimes this access involves the individual's control, sometimes not. The access model seems more complete than the control model. However, perhaps even this model is insufficient. Perhaps privacy cannot, on occasion, be defined with reference to access. To borrow an example from Adam D. Moore, imagine someone walking in a park. From this display, all manner of information can be garnered: the person's image, height, weight, eye colour and general physical abilities. Moreover, genetic material such as strands of hair will remain behind. Shed in public, this biological material is accessible to all. Hence the genetic data it contains, the very essence of that person's physiological identity, is accessible to all (Moore, 2003: 217-218).²⁹ If there is no restriction upon access whatsoever, does it not follow that all this genetic information is public, not private? If I make my DNA available by simply walking in public, do I not thereby renounce any privacy claims I might have to the information it contains? This seems illogical. Surely an expectation of privacy still prevails. Yes, access seems to be unlimited.

²⁹ Moore himself advocates the control model. However, he devises this thought experiment not to challenge the access model, but to show the inadequacy of Parent's definition of privacy as "the condition of not having undocumented personal knowledge about one possessed by others" (Moore, 2003: 217).

It is possible, albeit unlikely, that someone will find, analyse and share a DNA sample. However, it would seem that my right to privacy endures. This right would be violated only if the condition of privacy were breached, which would occur only if an unauthorised person were to perform a DNA analysis of my hair.

Certainly, with the advent of DNA testing technology, a new potential has arisen for our privacy to be violated in such a manner. This is why there exists a compelling urgency to adopt practical solutions in the form of legal and extra-legal measures, as I propose in chapter six. For now, though, my point is merely that Moore's hypothetical is *not* an effective counter-argument against the access model. The notion of restricting access remains definitive in this case. Even though access to my DNA is *potentially* unlimited on this scenario, it is only *actual* access that constitutes a breach of my condition of privacy and my right to privacy. What's more, that access would need to be unauthorised. As I will argue in chapter five, that authorisation depends upon two tiers of consent: individual consent; and the collective consent of the law. This two-tier model of consent, I argue, is usually definitive when it comes to internet privacy, given the extent to which digital technology allows for potentially unlimited access to our privacies, along the lines of Moore's hypothetical. Hence a further point is worth reiterating: as technology increasingly challenges the *condition* of privacy, the *right* to privacy is not necessarily affected. The right remains, even as the condition is breached. However, with the challenges brought by technology, the right to privacy can become unclear and contested, and hence must be articulated and protected in the clearest possible terms.

In light of the control/access debate, some scholars have sought to develop a hybrid. After surveying the classic legal and philosophical theories of privacy, Herman T. Tavani argues that several approaches provide important insights, but that none in itself provides an adequate account. Tavani then takes the best of the classic theories and incorporates them into one unified theory: the Restricted Access/Limited Control, or RALC, theory of privacy. This, he argues, can help us to frame an online privacy policy sufficiently comprehensive to cover the broad spectrum of privacy issues arising in regard to computers and information technology (Tavani, 2007: 1). Similarly, Moor uses the phrase "control/restricted access" to refer to his preferred model of privacy (Moor, 1997).

However, whereas privacy *always* involves a restriction upon access, privacy only *sometimes* involves control. As Ruth Gavison summarises, “... in its most suggestive sense, privacy is a limitation of others' access to an individual” (Gavison, 1980: 440). Privacy is invariably a matter of access. Sometimes an individual sets the limits; sometimes the limits are set by external forces. Sometimes privacy has to do with the intentions of the agent; sometimes it has to do with the intentions of others and society as a whole. In this vein, Jeffrey Reiman argues that the bedroom is an example of a space where privacy is about control. I decide who joins me in my bedroom. The toilet, by contrast, is a space where control is irrelevant; what matters is the mere fact that others are deprived of access. I don't decide; restrictions are set. As Reiman writes:

If we are to find the value of privacy generally, then it will have to be the value of this restriction of others. Sometimes its value will lie precisely in the fact that the restriction leaves room for our own control. But other times it will lie just in that others lack the access (Reiman, 2004: 198).

Hence I am adopting a wide conception of the access model, which allows for both control and externally-imposed restrictions. What's more, it is important to note that such a conception is substantively the same as a control/access hybrid. They are different in name only. On its own, the control model is too narrow. Only a wide conception of the access model, broad enough to accommodate control as well as externally-imposed limitations, can hope successfully to encompass the broad and complex sweep of privacy.

In this analysis, I have been seeking a conceptual, analytic account. Control is insufficient, but a broadly-defined conception of restricted access *is* sufficient. A thorough account of individual privacy acknowledges that privacy *sometimes* involves individually-imposed restrictions (that is, control) but that it *always* involves restrictions upon access (including by social norms and the law, as well as individual control, where appropriate). Similarly, a well-articulated hybrid of control and access is sufficient. We can thus come to an effective conceptual definition. Privacy involves a restriction upon access to ourselves, a restriction which sometimes involves control. Or, more fully:

The right to privacy is my right that others be deprived of unauthorised access to me and to information about me. In some cases, though not all, this right will involve my ability to control access to me and to information about me. The condition of privacy, meanwhile, is the state of others being denied that access.

The term “unauthorised” allows for both individual consent and the collective consent of the law, which are the subject of chapter five. Further, we might add to this definition, if we care to, that privacy is often connected with extra ingredients, including secrecy, intimacy and social relations. And we might add that different combinations of control, externally-imposed limitations and various other ingredients will come to the fore in different contexts. However, these extra ingredients are not necessary. Rather, what matters is that privacy is about access, and more specifically that I get to choose which swimmers I wear, but that my choice is circumscribed by external limits including social norms and the law.

Conclusion

As early as 1956, a federal judge in the US described privacy as a “hurricane in a haystack” (Schoeman, 1984: 200). Is privacy a condition? A right? If so, is it a descriptive right? A normative right? Or a legal right? And is it a value? An interest? A preference? Yes, privacy is all of these things. The concept is versatile and complex. We can have a private realm; we can distinguish bodily privacy from informational privacy; we can prise situational privacy from personal privacy. What’s more, further categories have been identified. In other words, privacy is widely inclusive. It is also contested. Hence our task of illuminating the ethics of internet privacy involves applying a contested notion in an ethically-contested realm. Still, conclusions can be drawn. In this chapter, I examined various accounts of privacy to show that the control model is inadequate, as revealed by my choice of bathing suit. Here is a choice about how much of my body I choose to expose to the world. I can be modest, exhibitionistic, or casually in between. However, the decision is not entirely mine. I cannot bathe nude. Privacy is about restricting access. Sometimes access is restricted by my choice; other times, access is restricted by externally-imposed limitations, such as the law, or social norms. In this way, I have shown that the best conception of privacy is the access model, widely-conceived, which equates to a control/access hybrid.

In *Rear Window*, Jeff is castigated for his voyeurism by his friend Doyle, a detective. “That’s a secret and private world you’re looking into out there,” Doyle

says. "People do a lot of things in private that they couldn't explain in public." Then, when the dog is found dead, its owner is distraught. "You don't know the meaning of the word 'neighbours'," she yells out into the courtyard. "Neighbours like each other, speak to each other, care if anybody lives or dies. But none of you do." After this outburst, the neighbours all run to their windows to see what's going on, except the salesman, who sits mutely in a darkened apartment, visible only by the glow of his cigar. In the next chapter, I turn to the question of why privacy matters, arguing that privacy is both non-instrumentally and instrumentally valuable for its connection to dignity, autonomy and relationships. Privacy is indispensable to us as individuals, and also as social beings. Proper respect for privacy must be balanced against a proper respect for publicity and community. As *Rear Window* shows, privacy is connected, somehow, to secrecy, and to relationships. What's more, it shows that privacy is sometimes about control, and sometimes about the law. On the one hand, residents can draw their curtains; on the other, the murderer is not entitled to keep his crime to himself. By the film's end, the killer has confessed, and various apartment dwellers have formed new relationships, seemingly with a healthier understanding of privacy, secrecy and social relations. In the closing scene, Jeff rests in his wheelchair, even more injured than at the start of the film. Lisa reclines nearby, reading a novel. As soon as Jeff falls asleep, Lisa puts down the book and opens a glossy magazine. Even from her boyfriend, she has secrets. The next question is, just why does she value these privacies?

Chapter 4

The value of relational privacy

Truman Burbank is completely unaware that his whole life is an elaborately-constructed fiction. A chirpy naïf, Truman thinks he resides in a coastal town. In fact, the 30-year-old lives under a giant dome on a Hollywood set, where every moment of his existence, right from birth, has been filmed as part of a reality TV show. Complete with product placement and choreographed extras, Truman's life is broadcast live to an audience of billions. Of course, Truman Burbank is not real; he's the protagonist of Peter Weir's 1998 comedy/drama *The Truman Show*, starring Jim Carrey. The drama raised the question of what life might be like for an individual whose privacy has been taken away entirely, not by a government, but by a corporation. Slowly, as the truth begins to dawn, Truman's quest becomes to escape his faux reality so that he might find his way into a more human world, where privacy (among other things) is still possible. As Truman begins to suspect something is amiss, his millions of viewers become even more enthralled. Will he discover the truth? How will he react? Just how important is his privacy?

In this chapter, I address a simple question: Why does privacy matter? Again, the question's simplicity belies the complexity of proposed answers. In section one, I propose two worlds: one without privacy; one with absolute privacy. Both are inhospitable, uninhabitable and, indeed, impossible. Hence I propose a privacy axis, on which every society can be charted according to how much it values privacy. In sections two, three and four, I explore in detail three justifications for the value of privacy: dignity; autonomy; and relationships. Each of these, I will show, are significant. Contra the philosophers who argue that humanity would be better served with full publicity, I argue that privacy is a fundamental good. The condition of privacy matters, and the right to privacy is worth protecting, though for different reasons at different times. In some cases, privacy matters for reasons of dignity; in others, it matters for reasons of autonomy, and particularly relational autonomy; in still others, it matters for our relationships. Often, it matters for some combination of the above. I then sketch an outline of *relational privacy*, which recognises that individuals are only ever beings-in-relation, and which

acknowledges that privacy is also a social and public good. It is fundamentally valuable for individuals, but also for society and democracy. Privacy is a means for individuals to separate and withdraw, certainly, but also a means for them to bond and coalesce. The nature of privacy's value will shift according to circumstance, and that invariably privacy must be weighed up against competing rights and freedoms. On the privacy axis, societies and individuals ought to strive to situate themselves somewhere in the centre, abiding by ethical and legal prescriptions that affirm privacy's fundamental, albeit not all-conquering, value.

I – The privacy axis

Over breakfast one morning, shortly after turning five, my daughter Lola looked pensive. "I'm thinking in my head," she said, finally. "Oh," I responded, between slurps of coffee. "What are you thinking?" She turned to me and said, coolly, "That's a secret. A secret only for me." Having barely started kindergarten, my daughter already had a sense of some private part of herself, a part that she could choose not to share. In the previous chapter, I argued that privacy involves restrictions on access. Clearly my daughter agrees. Moreover, I argued that there is sometimes a link between privacy and control. It seems Lola perceives such a link. Our exchange also reveals another significant point: as long as we can have our own thoughts, we have a degree of privacy. That is, if others cannot read our minds, we have a degree of the *condition* of privacy. As such, it seems that a world without privacy is impossible (at least currently).³⁰ The nature of human existence is such that each individual has his or her separate consciousness,

³⁰ Which is not to say that mind-reading is theoretically impossible. Perhaps the internet will make it possible. Perhaps individuals will one day be able to transcend their own consciousness, potentially in a type of singularity, as described in chapter one. Indeed, perhaps the internet is already beginning to make it possible to read users' thoughts. We have seen (in chapter one) the way location data can be used to predict a user's location 24 hours into the future. We have also seen (in chapter two) the way Facebook likes can unintentionally reveal supposedly hidden personality traits including sexuality and political views. Meanwhile, "emotion recognition" software is being used by advertisers to gauge users' latent emotions by analysing facial cues. The goal is to create a "mood-aware" internet that reads a user's emotions to help then shape their content (Matheson, 2014). By deducing facts that users don't even know they're sharing, these developments can be arguably regarded as types of mind-reading. What's more, my argument assumes that our thoughts are indeed our own, and not the result of manipulative techniques that potentially compromise the autonomy of users' thinking. The phrase "filter bubble" captures the way content is manipulated and personalised for each user on the net (Pariser, 2011: 1-3). Emotions can be manipulated too, with research showing that emotional states can be transferred between Facebook users via "emotional contagion" (Kramer et al., 2014).

thereby guaranteeing some degree of privacy. A person can say one thing while thinking something contradictory. In the absence of mind-reading, in the absence of the ability to transcend individual consciousness, every individual has privacy. “L’enfer, c’est les autres,” wrote Jean-Paul Sartre in 1944, in a decidedly pessimistic account of the unknowability of another’s thoughts.³¹ My daughter was more upbeat about the unknowability of her thoughts. For her, this small privacy was something precious.

However, some philosophers argue that privacy has no value at all. They propose that people and societies could, and *should*, exist without it (see Schoeman, 1984: 200). Privacy exists, they argue, merely due to the illusion that elements of one’s life are embarrassing and unique; the reality, however, is that humans’ lives are essentially universal, and that privacy-related embarrassment and uniqueness are feelings that, with sufficient progress, we could and should discard. For his part, Plato considered privacy as an obstacle to the ideal state. In *The Laws*, Plato argued that a state will be truest, best and most exalted in virtue if it contains:

... this communion of women and children and of property, in which the private and individual is altogether banished from life, and things which are by nature private, such as eyes and ears and hands, have become common, and in some way see and hear and act in common, and all men express praise and blame and feel joy and sorrow on the same occasions, and whatever laws there are unite the city to the utmost ... (The Laws, ch. 5, §738d-e, quoted in Moore, 2013b: 3)

Plato’s vision, not unlike Asimov’s vision of Gaia (see introduction), has modern supporters. As Schoeman writes, “People who hold this view claim that institutions of privacy are conducive to social hypocrisy, interpersonal exploitation through deception, and even a-social or anti-social loyalties” (Schoeman, 1984: 200). On this view, privacy is atavistic and selfish. A world without privacy would be a world of brilliant sunshine. It would be a world without hypocrisy, shame and deception.

Who is right? Lola or Plato? A preliminary point is that privacy norms vary dramatically over time, and across cultures. An illustration lies in the ruins of the ancient Greek city of Ephesus, in modern-day Turkey. There, modern tourists can take their position on one of a series of toilet seats in the public hall, where posh

³¹ The line translates as, “Hell is other people” (Sartre, 2005). Intriguingly, 2013 saw the launch of Hell Is Other People, billed as a social media app for people who hate social media, which uses location-based information to allow users to avoid those they don’t like (Bosker, 2013).

Ephesians “gathered to commune, two thousand years ago, as they collectively emptied their bowels” (Whitman, 2004: 1154). The variation is clear too in modern social norms: in the US but not Europe, people casually discuss salary and net worth; in Europe but not the US, people casually take off their clothes (Whitman, 2004: 1158). To understand privacy, anthropologists have charted the wildly diverging norms and customs that prevail among Native Americans, Polynesians and Javanese, among many others (Moore, 2013a: 221-222). Does this suggest that privacy is optional? Quite the reverse. Despite the divergences, there appears to be one constant: in one form or another, privacy is valued, recognised and institutionalised in all human societies (Murphy, 1964: 1257; Westin, 1967: 12). Privacy is a cultural universal, it seems, necessary for humans to survive as the social animals they are (Moore, 2013a: 222).³² Privacy exists among the Tuareg and the Thinglet, among Australians, Americans and Europeans, albeit in wildly varying forms. Of course, this is hardly proof that privacy *should* be valued.

Let’s propose that every society can be charted on some sort of axis, designed to represent the extent to which a society values privacy. At one end of the axis is a society that values only that which is public, and that dismisses privacy as worthless and irrelevant; in this society, privacy is not valued at all. At the other end of the axis is a society that values privacy fiercely, just as much as is possible. In other words, here privacy is as close to absolute as possible. This axis, then, represents how much a society values privacy by taking account of several factors, including: the extent to which individuals enjoy the condition of privacy; the extent to which the right to privacy is protected in law; the extent to which social norms favour privacy; and so on. It might look something like this:

No Privacy

Some Privacy

Absolute Privacy

³² The value of privacy may not be exclusive to humans. In 1967, Alan Westin wrote, “virtually all animals seek periods of individual seclusion or small group intimacy”; this is “territoriality, in which an organism lays private claim to an area of land, water, or air and defends it against intrusion by members of its own species” (Westin, 1967: 8).

Granted, such an axis is reductive, given the complex concept that we have established privacy to be. Still, the task of charting any given society on such an axis is theoretically possible. We have already seen, for instance, that Europe's privacy laws are considerably more stringent than those of the US or Australia. Based on these and other measures, Continental countries would be plotted nearer "absolute privacy" than the US and Australia, which would be plotted nearer "no privacy", given their high regard for competing considerations such as freedom of expression, the free flow of information, the government's duty to act in the interests of national security, and so on. What's more, we have heard suggestions from Mark Zuckerberg and others that people's privacy norms are shifting and loosening in countries such as the US, suggesting that some societies may be moving closer to "no privacy".

Let us then contemplate two imaginary societies: one which places no value on privacy whatsoever, and in which privacy's condition and right have thereby been eradicated; and another which places an absolute value on privacy, and which seeks to privilege both its condition and right over any competing considerations. What might these societies look like? What would it be like to live there?

In the first, everything is public that could possibly be public. There are neither doors nor walls on toilets; there are no curtains on bedroom windows. There is no ability to control the flow of information, even if sensitive; and there is no restriction upon access to oneself or to information about oneself. Every space is public, all data is public, every body is public. This is the total lack of privacy contained in the notion of an everpresent, omniscient god (Gavison, 1980: 443). Or in the world inhabited by Truman Burbank, who is watched over by the godlike TV producer Christof. At the other extremity, privacy is worshipped as a supreme value. Every house has its curtains drawn; every property has a fence; and privacy-protecting laws have been drafted to prescribe that one individual cannot take a photo of another, even with consent, and that every person must be fully covered, at all times. Spaces are private; data is private; bodies are private. Secrecy, anonymity and solitude abound. This world, I suggest, is even harder to imagine than the first.

There is something deeply troubling about both these worlds. In the first, the bright light of full publicity is blinding. Each detail, intimacy and body is exposed. Everything is shared and social, with nothing left for individuals (or couples, or families) to keep to themselves, and to share judiciously. In the second, the light cannot penetrate. Dark and shrouded, this is a place of secrets and shadows. Here, the individual has been isolated at the expense of interaction, community and social engagement. In both worlds, neither individual nor society could possibly flourish. Rather, these two scenarios are both, in completely contrasting ways, dystopian. As Gavison writes, “We start from the obvious fact that both perfect privacy and total loss of privacy are undesirable ... Privacy thus cannot be said to be a value in the sense that the more people have of it, the better (Gavison, 1980: 440). Fortunately, both worlds are also impossible, as suggested by the difficulty of mapping their contours. As Gavison also writes, “the total loss of privacy is as impossible as perfect privacy” (Gavison, 1980: 428).

By contemplating societies at either extremity of a privacy axis, we arrive at the strong intuition that too little privacy is undesirable, and too much privacy is undesirable too. Privacy ought to be valued, but not *in extremis*. Countervailing rights and freedoms must be valued too. In this way, it would seem a balance must be struck to recognise that my right to privacy is merely one of several rights, which must in turn be balanced against your equal rights (Fried, 1968: 478). After co-writing the 1890 essay “The right to privacy”, Louis Brandeis reportedly planned to write a companion piece on “The duty of publicity” (Berger, 2009). He never published such a piece, but did pen the oft-quoted lines: “Publicity is justly commended as a remedy for social and industrial diseases. Sunlight is said to be the best of disinfectants; electric light the most efficient policeman” (Brandeis, 2009: 62). For individuals and society to prosper, it seems a balance must be struck between the light of publicity and the shade of privacy. To help strike such a balance, I turn now to elaborating more precisely just why and when privacy ought to be valued.

II – Dignity

Scholars often invoke several reasons when arguing for the value of privacy, and these reasons are sometimes difficult to disentangle. For instance, Charles Fried

writes of privacy's connection to respect, individual integrity and social personality, which, on his view, are all linked (Fried, 1968). And Stanley Benn, seeking to apply the overarching notion of respect for persons, invokes both dignity and autonomy when he writes of privacy's role in enabling individuals to be self-aware subjects able to pursue their own projects (Benn, 1971). In what follows, I will (as far as possible) disentangle several strands, beginning with an argument based on human dignity and respect for persons. Ultimately, however, I too will be arguing that privacy matters for several reasons, and that different combinations of these reasons come to the fore in different circumstances. As with Fried and Benn, many of my arguments have Kantian underpinnings.

Many modern justifications for privacy can be traced back to the 1890 essay "The right to privacy". By arguing that every individual deserves respect on account of his or her "inviolable personality", Warren and Brandeis opened the way for subsequent justifications of privacy based on dignity and respect (Warren and Brandeis, 1890: 205, 211; see chapter three). In 1964, Edward Bloustein argued that the notion of "inviolable personality" encompasses notions of individual dignity and integrity, personal uniqueness and personal autonomy. "I take the principle of 'inviolable personality' to posit the individual's independence, dignity and integrity; it defines man's essence as a unique and self-determining being" (Bloustein, 1964: 971). It is our concern for these values, Bloustein argued, that unifies our conception of privacy; and intrusion upon one's seclusion is not just a threat to emotional tranquility, but also an affront to human dignity. Four years later, Charles Fried also advanced a Kantian account founded on respect (Fried, 1968). For both Bloustein and Fried, privacy is so intertwined with dignity that an attack on privacy will be an assault on a person's very personhood. Privacy is linked to dignity and individuality, they argue, and an attack upon privacy is potentially an attack upon what it means to be human.

Following Bloustein and Fried, I am claiming that some intrusions upon privacy are violations of dignity. And by dignity, I use the term in a Kantian sense to mean that which marks out the priceless worth of humanity, and which thus demands respect (Kant, 2009: 434-435; Wood, 1999: 140; Formosa, 2017, in press: introduction). That is, I mean status (or inviolable) dignity, rather than the achievement (or aspirational) dignity that we can earn for ourselves through our

actions, appearance, reputation, and so on (Formosa, 2017, in press: introduction). In the *Groundwork*, Kant wrote that “if [something] is exalted above all price and so admits of no equivalent, then it has a dignity ... [and] morality, and humanity so far as it is capable of morality, is the only thing which has dignity” (Kant, 2009: 434-435; see further discussion in chapter five). The correlative of dignity is respect: for Kant and Kantians, it is human dignity that entitles every person to respect (Fried, 1968: 479; Wood, 1999: 140). What’s more, to make the claim that privacy can be justified on grounds of dignity is to cite its non-instrumental value (Fried, 1968: 477).³³ This is a contested point. In 1960, legal scholar William Prosser argued that privacy really doesn’t matter so much after all. For Prosser, privacy has only instrumental value, and privacy violations are founded more specifically on injured reputation and emotional distress.

Can we test whether privacy, via its links to dignity, has non-instrumental value? Let’s say I have 100 Facebook friends, with whom I share, in a carefully discerning way, personal anecdotes, confessions and photos. Let’s further suppose that third parties, including companies and government agencies, have access to this personal material without my knowledge. Further, let’s suppose that these companies and agencies can make many other highly revealing deductions about me from other sources, including from the company I keep and my offline behaviour. Instead of merely sharing select privacies with limited friends, as I believe I am doing, I am in fact sharing these privacies and many more besides with a wide range of friends, strangers, companies and agencies. Even if there has been no instrumental consequence to this encroachment upon my privacy, there has, it seems, been a failure to respect my dignity. At a fundamental level, my humanity has not been respected. Indeed, arguments contra Prosser have tended to cite cases in which a person’s privacy is violated without that person knowing of the violation. Hence James T. Moor posited an unseen voyeur, Tom, who uses secret cameras and other devices to record everything about someone. To suit our purposes, let’s imagine Tom as the webcam hacker who spied on Cassidy Wolf (see chapter two). But unlike that hacker, let’s imagine that Tom doesn’t do her any direct harm. He doesn’t share his footage with anyone; he doesn’t harass

³³ Fried uses the term “intrinsic” rather than “non-instrumental”, as does Moor. The terms are sometimes used interchangeably. However, I prefer “non-instrumental”, which here denotes something valuable for its own sake. Intrinsic denotes something whose value is self-generated. The contrast will become clearer below, when I draw on Reiman’s distinction between the intrinsic and extrinsic value of privacy (Reiman, 2004).

Wolf; and he doesn't attempt blackmail. The only benefit he gains is the fulfilment of his own voyeuristic impulses, of which the object of his voyeurism is completely unaware. Is Tom harming this woman? I would argue yes, just as Moor does: "I think most of us will agree that there is something repugnant about Tom's peeping ... Some people, including myself, regard privacy as intrinsically valuable, not merely instrumentally valuable" (Moor, 1997: 28-29). The harm here, it seems, comprises an attack on dignity. By exploiting Wolf, Tom is compromising her humanity. If Wolf knew, we can only imagine she would try to stop Tom's peeping, just as my Facebook use has changed now that I know that my privacies are being widely shared.³⁴

As we have seen in previous chapters, unseen voyeurism is commonplace on the net. The Creepshots website, for which men secretly take and share photos of women's buttocks and breasts, defends itself with the motto, "No harm, no foul" (@CreepShot, 2017). The suggestion is that because these women will never know that men have photographed them, they will be unharmed when others engage in mutual, secret voyeurism. However, there is harm. As women's privacy is violated, dignity is compromised. Without their knowledge, their bodies have been objectified and exploited, used merely as a means to satisfy the desires of others. And with the compromise of their dignity something larger has occurred: the worth of humanity as a whole has been diminished. Prosser argued that a privacy violation required injured reputation or emotional distress. His justification is too narrow, particularly in a digital age. For 30 years, Truman Burbank was unaware that his every breathing moment was secretly being watched by an audience of millions. Such watching may not have any instrumental effects; non-instrumentally, however, the violation is egregious.³⁵

³⁴ As detailed in chapter two, the real-life Cassidy Wolf did go to the police to stop the spying (although the real "Tom" also engaged in blackmail and harassment). What's more, my social media use continues to change: the more I learn about the threats to privacy online, the more I self-censor. Am I typical? The jury is out. Researchers have found that users understand the risks and care about privacy on the internet, but then act in a way that contradicts that care (Debatin et al., 2009; Taddicken, 2014). However, researchers have also shown that many users are changing their online behaviour given their increasing awareness of privacy risks (Peterson, 2016).

³⁵ If privacy is linked to dignity, and dignity ought not be for sale, then there is an argument to be made that privacy, or at least parts thereof, ought not be for sale. That is, that privacy, or at least those portions of privacy connected with dignity, should not be able to be bought and sold. Without addressing privacy specifically, communitarian Michael J. Sandel argues that the market economy has given way to the market society, where an increasing commodification of life is leading to greater inequality. The rich can now, for instance, buy their children places in academically elite schools. The market, Sandel argues, inevitably changes the character of the goods it touches, such as sex, friendship, family life, health, education, art and more. This is

Not all dignity-based arguments are based on cases where the privacy violation is surreptitious. Another strand of arguments invokes cases where people are at their most vulnerable, such as when women are giving birth. In the labour ward, a woman generally has an expectation that she will be surrounded only by designated intimates, as well as requisite medical professionals, but not utter strangers and random voyeurs. A failure to respect a woman's right to privacy during childbirth, it seems, debases her dignity. As Bloustein writes: "A woman's legal right to bear children without unwanted onlookers does not turn on the desire to protect her emotional equanimity, but rather on a desire to enhance her individuality and human dignity" (Bloustein, 1964: 973, 982). Childbirth is universal. Everyone has attended at least one. Yet it can also be a powerfully private event. Further, it is a potent illustration of the way privacy can simultaneously have both instrumental and non-instrumental value. If an unwanted intrusion upon childbirth causes ongoing emotional distress, there are clearly potential instrumental ramifications, starting with a stalled labour, but there has also been a non-instrumental failure to respect the woman's humanity and to exercise the virtue of compassion. Respecting humanity and exercising compassion are two principles we ought to follow for their own sakes, for the non-instrumental imperative of furthering the end of humanity (to use Kantian terms). The instrumental benefits are welcome too, of course, but we respect humanity and exercise compassion even when no instrumental good will ensue. Indeed, although I agree with Bloustein's invocation of dignity in this case, I propose that "emotional equanimity" *should* be relevant. Emotional distress is precisely what we hope to avoid, for reasons that are both instrumental and non-instrumental. Bloustein himself seemingly recognised this point when, in giving his account of the US case law on privacy as it stood in 1964, he wrote of the

having a deep effect on democracy, and hampering citizens' ability to share in a common life (Sandel, 2012: 202-203). As I have shown, privacy is being bought and sold, by companies including Acxiom and Facebook. Indeed, the commodification of privacy is already so well-established on today's internet that to rail against it smacks of King Canute. Still, does selling privacy fail properly to respect human dignity? Are we headed for a world where the rich are able to afford privacy (via superior software, hardware and IT knowledge), but the poor are not? And if we as citizens decide that we should allow privacy to be bought and sold, we need to ask: to what extent? Sandel writes, "We need to think through the moral limits of markets. We need to ask whether there are some things money should not buy" (Sandel, 2012: 7). These issues are beyond the scope of this thesis, but warrant attention. It is also interesting to note that in *The Truman Show* Truman Burbank was the first baby legally adopted by a corporation. The film's writers presumably saw a strong potential link between a total absence of privacy and a dignity-compromising "ownership" of humans by companies.

“spiritual characteristic” of privacy cases. Yes, he wrote, sometimes privacy has instrumental value, but it is true also that “the interest served in the privacy cases is in some sense a spiritual interest rather than an interest in property or reputation” (Bloustein, 1964: 1002).

These cases show that some invasions of privacy amount to violations of dignity. In cases of secret surveillance such as Creepshots, women (and men) have been treated as objects, as *things* whose feelings, reason and worth are irrelevant, and certainly not worthy of respect. The violations constitute a violation of the priceless worth of humanity, both as it pertains to the specific victims of self-styled “creeps”, but also as it pertains to humanity generally. The story is similar when the privacy of the vulnerable is violated. When a woman’s privacy is violated during childbirth, there may be direct consequences. Her labour may well become complicated. However, there is also something more, in the way the woman herself, and in the way humanity more widely, has not been afforded due respect, and has been treated without compassion. In the ultimate life-affirming moment that is birth, a failure to respect privacy is a particularly significant and symbolic violation of dignity. However, it is easy to imagine many other privacy invasions of the vulnerable (children, the mentally ill, the elderly) that violate dignity. Admittedly, privacy-justifying arguments based exclusively on dignity are hard to prosecute. Dignity is difficult to define and difficult to defend. There are at least five different conceptions of dignity (Formosa, 2017, in press: introduction). I have been invoking Kantian dignity, or status dignity. Even this conception has its detractors, with Stephen Pinker attacking the “stupidity of dignity” and Ruth Macklin dubbing it a “useless concept” (see Formosa, 2017, in press: introduction). No wonder that scholars who justify privacy on the grounds of dignity usually do so in conjunction with justifications based on autonomy, relationships, and more. There may be occasions when a privacy violation is most egregious simply and precisely because it violates status dignity. Perhaps the social media and Peeping Tom examples outlined above are two of them. Perhaps childbirth is another. Ultimately, to justify privacy violations on the grounds of dignity alone is to submit to an unnecessary constraint, as we will see when we turn our attention to the justification from autonomy.

III – Autonomy

The aim of the previous section was to show that privacy matters because some violations of privacy amount to violations of status dignity. If we can protect people from secret surveillance, and if we can prevent intrusions on vulnerable people, then in some cases it will be dignity that we are protecting. In this way, I have been arguing that privacy is an individual good. As I turn to the closely related justification from autonomy, I will be arguing that privacy is also a social good. Simply, I will argue that some violations of privacy have the potential to compromise autonomy. More specifically, I will argue that some privacy encroachments restrict our ability to think freely, act freely, test out unpopular opinions and express ourselves fully. These encroachments may adversely affect our moral integrity, promoting conformism and adherence to what is conventionally regarded as right, rather than a reasoning and self-legislated adherence to what is *right*. As Vallor notes: “Surveillance technologies that work too well in making us act ‘rightly’ in the short term may shortchange our moral and cultural growth in the long term” (Vallor, 2016: 191). First, I will give an account of individual autonomy and its connection to privacy. I will then explicate a relational conception of autonomy, before developing the concept of relational privacy.

The notion of individual autonomy has come to occupy a central place in moral and political philosophy. For Kant, autonomy played an unconditional role in human morality, such that individuals should act in such a way as to respect and promote the ability of others to determine their own lives. Indeed, Kant regarded autonomy as the foundation of human dignity (Kant, 2009: 436). Dignity involves observing the moral law, but also being autonomous with respect to it by self-legislating one’s morality, and for Kant the antithesis of an autonomous will is the heteronomous will, which is governed by something other than a self-given law of reason (Kant, 2009: 444). However, there is debate about what precisely Kant meant by autonomy. Formosa’s interpretation is that Kant’s conception of autonomy amounts to a unified theory of moral and personal autonomy, but this is contested (Formosa, 2013b).

Today, the value of individual autonomy is rarely questioned in applied ethics and legal philosophy (Mackenzie and Stoljar, 2000: 4). However, just as there are varying interpretations of Kant's use of "autonomy", there remains little consensus as to what the concept means: bioethicists often construe autonomy as informed consent; liberal political theorists consider autonomy in terms of the right to non-interference; and Rawlsian liberals tend to think of autonomy in Kantian terms as the capacity for rational self-legislation (Mackenzie and Stoljar, 2000: 5). One suggestion is that while various theorists all invoke the concept of autonomy, they are invoking different conceptions which derive from different normative frameworks, value orientations and political commitments (Mackenzie, 2014: 15). Catriona Mackenzie responds by arguing that, "autonomy is a multidimensional, rather than unitary, concept". Mackenzie proposes a taxonomy "that distinguishes three distinct, but causally interdependent, dimensions or axes of autonomy: self-determination, self-governance, and self-authorization" (Mackenzie, 2014: 15-16). The self-determination axis identifies external, structural conditions for autonomy; the self-governance axis identifies internal conditions, which comprise having the requisite skills and capacities to choose and to act; and the self-authorisation axis involves regarding oneself as authorised to exercise practical control over one's life (Mackenzie, 2014: 17-18). Together, these three causally interdependent axes comprise the defining characteristic of free moral agents. In this section, I follow Mackenzie to define individual autonomy not as a specific conception, but as a broad, multi-dimensional concept. It is, in lay terms, the ability to steer one's own ship.

Unsurprisingly, autonomy figures prominently in the privacy literature. Bloustein (1964), Fried (1968) and Benn (1971) all developed explicitly Kantian justifications for privacy involving autonomy. Joseph Kupfer argues that "privacy is essential to the development and maintenance of an autonomous self" (Kupfer, 1987: 82); and Beate Röessler argues that privacy is a necessary condition for individual autonomy (Röessler, 2005: 42-76). As Adam D. Moore summarises: "According to these theorists, privacy is morally valuable because it protects and promotes the sovereign and autonomous actions of individuals – since autonomy is morally valuable privacy must be as well" (Moore, 2013b: 10). In its strongest form, the argument runs that there can be no autonomy without privacy. This claim seems, intuitively, too strong. Not every invasion of privacy, I propose, will

have an impact on autonomy. If I spend a few minutes looking through my 11-year-old daughter's diary while she's at school, this may compromise her capacity to be self-determining, as it affects the external, structural conditions of her actions in the context our relationship. She doesn't know I have seen her entries, and so when we next interact she will be unaware that I know something private about her. My intrusion is an intrusion on her freedom. If, however, I look at something more trivial, such as the playlist of songs she has created, her autonomy may be undisturbed. As Moor argues, privacy and autonomy come apart: "Privacy is not an essential condition for autonomy. It is conceivable to have autonomy without privacy" (Moor, 1997: 29). Some privacy violations, particularly trivial ones, will have no impact on autonomy.

A second version of the argument is that people tend to act differently if they think they are being watched. Famously, US politician Hubert Humphrey said: "If we can never be sure whether or not we are being watched and listened to, all our actions will be altered and our very character will change" (quoted in Reiman, 2004: 194). This version immediately invites objections: first, it does not follow that people's autonomy is compromised just because they are behaving differently; second, people are often being watched in a way that seemingly does not compromise their autonomy; third, people's privacy has not been violated just because they think they are being watched. This version of the argument is often invoked, but lacks rigour. A third version of the argument is less stringent, proposing simply that *some* invasions of privacy compromise autonomy. This is the claim I will be defending in this section: that harming privacy *may* compromise autonomy.

In relation to the internet, autonomy is a commonly-cited justification for privacy. I have previously mentioned Creepshots (above) and Google Glass (chapter two). Ironically, on its website, Creepshots says the "creeps" who surreptitiously take and share photos ought, if challenged, to invoke their right to privacy: "If you see someone trying to catch you by looking over your shoulder at your phone/camera then politely tell him/her to stop invading your privacy" (CreepShots, 2017). Clearly the "creeps" consider that such invasions of privacy infringe their autonomy. Meanwhile, Google Glass arguably failed because of user privacy concerns. As the online magazine *Digerati* noted:

Google Glass created an environment where people were subjected to the potential for ‘always on’ recording. Living like this, with a constant fear of being caught on camera, alters how people behave. Glass also gave prospective stalkers and creeps in general the ultimate tool for taking invasive photos of women in public without their knowledge (Edwards, 2016).

The author thus proposes both dignity (“invasive photos”) and autonomy (“alters how people behave”) as justifications for privacy. After disparaging Glass, the author then praises another set of internet-connected glasses, Specs by Snap Inc., which clearly signal when the wearer is filming or taking a photo (Edwards, 2016). Again, we need to take care to distinguish varying claims. Here, it seems, the author is suggesting that the prospect of secret surveillance by Google Glass is wrong because it alters behaviour, whereas Specs by Snap are acceptable because they signal when recording is occurring. The author also implies, but doesn’t state, that lack of consent suggests a violation of privacy and a compromise of autonomy. With Google Glass, there can be no consent because people can never know if they are being recorded; however, Specs by Snap offer the possibility of consent by letting people know if they are being recorded. Clearly, Specs by Snap are preferable to Google Glass. Does this mean they get an ethical thumbs up? Not necessarily.

I now want to focus on three ways in which privacy invasions can limit autonomy: first, the way invasions of privacy can inhibit the “self-creative enterprise” of living our lives; second, the way invasions of privacy can curtail alternative viewpoints and promote conformism, thereby creating the preconditions for totalitarianism; and third, the way invasions of privacy can stifle our ability to be moral agents.

For Stanley Benn, invasions of privacy constrain the pursuit of the self-creative enterprise of living. Benn argued that people act differently when they are among others. When we are alone, we act a certain way; when we are among others, we become conscious that we are being judged from the others’ perspectives, and hence we are liable to act differently (Benn, 1971). Indeed, this is so if we merely *suspect* we are being observed, even if in fact we are not. Benn thus prescribed that we should not watch others against their will, unless there are strong reasons to do so. For Benn, we ought to realise that others have a point of view on the basis of which they make choices, and we ought to respect those choices, unless

there are compelling countervailing reasons. Hence Benn argued against clandestine surveillance. Even when secret surveillance does not affect the behaviour of the person being watched, it does undermine that person's ability to make rational, well-informed choices, simply from the fact of not knowing about the surveillance. My choices will not be fully-informed if I am under surveillance and do not know about it. If I am being watched without knowing it, I will act a certain way; if I were to know of the surveillance, I might act differently. Secret surveillance knowingly and deliberately alters the conditions of the person being spied upon, and thus fails to respect them as a person.

Respect for someone as a person, as a chooser, implies respect for him as one engaged on a kind of self-creative enterprise, which could be disrupted, distorted, or frustrated even by so limited an intrusion as watching (Benn, 1971: 26).

The concern for Benn is that watching, and the trespass upon privacy that such watching entails, can compromise an individual's autonomy. This claim is easy to justify if the surveillance is known. Yet even if the surveillance is secret, autonomy suffers. Had I known I was under surveillance, I might have acted differently. Anyone engaging in such secret surveillance is failing to express respect for me. What's more, my autonomy is being undermined to the extent I am no longer able to make informed choices. Self-determination is crimped. The self-creative enterprise of living is stunted.

Above, I raised the example of peeping Tom, who is spying on Cassidy extensively, without Cassidy knowing. I posited that Tom neither shares nor exploits this footage (aside from satisfying his immediate voyeuristic urges). I argued there that her dignity had been violated. To this I would add that her autonomy has suffered too, along the lines suggested by Benn. Perhaps her behaviour would change were she to know of the watching. In any case, she is conducting her life without full knowledge of the prevailing circumstances. Further, what if we consider autonomy in light of what is happening *globally*? What if on a wider level Cassidy and other users know that this sort of spying is not uncommon on the internet, with hackers, companies and government agencies all able to surreptitiously access webcams? Even if Cassidy doesn't know she herself is being spied upon, she may well know that such surveillance is a distinct possibility, and this may curtail her ability to act as she would like. In response, users might put sticky tape over their webcams, like Mark Zuckerberg, or try

other DIY counter-surveillance techniques. Or they might modify their behaviour in a way they would prefer not to, but feel is necessary given the knowledge that their smartphone, laptop or internet-connected fridge might be enabling covert surveillance. Tom's surveillance is, it seems, having an effect on Cassidy's autonomy, and on autonomy globally. In Cassidy's case, the effects on autonomy are subtle and indirect. In other cases, invasions of privacy can *clearly and directly* compromise autonomy. If my car is being tracked, I may not drive places I would otherwise wish to go; if my every click is being recorded, I may not visit websites I would otherwise frequent.

Philosophers have also argued that invasions of privacy can lead to conformism. In *On Liberty*, Mill wrote that liberty is a school for character: "A person whose desires and impulses are his own - are the expression of his own nature, as it has been developed and modified by his own culture - is said to have a character" (Mill, 2011: 112). To this, Reiman adds that *privacy* is a school for character, sheltering people from conformity and allowing them to become the sorts of people who are not vulnerable. If our goal is to foster strong-willed citizens who are able to resist social pressures, then we must first give them privacy, in order that they can gain experience making and acting upon their own judgments. While they are vulnerable, they need privacy, in order to become the sorts of people who are less vulnerable to conformity. "In short, the vast majority of actual people need privacy for free action, and those who do not, needed privacy to become that way" (Reiman, 1976: 203). Meanwhile, Gavison writes that in certain spheres of life, including artistic expression and intellectual development, people need freedom from close and constant scrutiny to flourish. It is privacy that affords people the space, both intellectual and emotional, to contemplate unpopular ideas without the pressure of social disapproval and sanctions. "Privacy is needed to enable the individual to deliberate and establish his opinions" (Gavison, 1980: 450). For Gavison, privacy can act as a shield. If a person is gay, then homosexuality is their individual standard, but this contradicts the social norm of heterosexuality. Even in a liberal society, beliefs or behaviours that stray from the norm can arouse hostility. This prospect may thus inhibit a person from engaging in a homosexual relationship. It is privacy, Gavison argues, that ensures social norms do not govern such behaviour (Gavison, 1980: 452-453). This is particularly valuable when social norms are, say, homophobic, racist or

misogynistic. For Gavison, privacy can safeguard against conformism (Gavison, 1980: 463-464).

For Bloustein, similarly, the loss of autonomy attending invasions of privacy render a person purely conventional. This loss will lead her to abandon her individuality and become part of an amorphous mass through her desire to conform to others' expectations.

The man who is compelled to live every minute of his life among others and whose every need, thought, desire, fancy or gratification is subject to public scrutiny, has been deprived of his individuality and human dignity. Such an individual merges with the mass. His opinions, being public, tend never to be different; his aspirations, being known, tend always to be conventionally accepted ones; his feelings, being openly exhibited, tend to lose their quality of unique personal warmth and to become the feelings of every man. Such a being, although sentient, is fungible; he is not an individual (Bloustein, 1964: 1003).

Bloustein's paper is a lengthy response to Prosser's argument that privacy has only instrumental value. For Bloustein as for Prosser, privacy certainly does have instrumental value. It can protect reputations and prevent emotional distress. However, for Bloustein privacy also has non-instrumental value, including in the way it fosters autonomy. If a woman's privacy is violated, her freedom to act as she chooses and to think as she chooses may be compromised. Her acts and beliefs may come to resemble those conventionally accepted, and her opinions may align more closely with the majority view. This, of course, is not certain. Some argue that privacy is a necessary condition for autonomy. My claim is less strong: that *some* invasions of privacy involve a violation of autonomy. Certainly, privacy is not a sufficient condition for autonomy. Just because my privacy is not violated, does not mean I am acting autonomously. My suggestion here is simply that invasions of privacy tend to lead to conformism, which in turn suggests a curtailment of autonomy.

Given the activities of hackers, Acxiom and the Australian Signals Directorate (see chapter two), internet users can fairly suspect that every click, drag and interaction is potentially being observed. In the context of such unwanted surveillance, a user may alter her behaviour. She may not visit certain websites; she may not write certain emails; she may not post certain photos. We have already seen that many people are behaving differently online because of privacy concerns: they are filtering, if not halting, their posts to social networks; they are

no longer buying goods online; and they have stopped expressing their opinions in comment threads and forums (Peterson, 2016). Further, a 2016 study revealed that perceptions of surveillance practices can stifle the expression of minority political views. This, in turn, can have a chilling effect on democratic discourse. “Knowing one’s online activities are subject to government interception and believing these surveillance practices are necessary for national security play important roles in influencing conformist behavior” (Stoycheff, 2016: 297).

Jeffrey Reiman argues that losses of privacy can lead to the risk of “psychopolitical metamorphosis”. Here is the idea that a lack of privacy infantilises people, by impoverishing their inner life and making them vulnerable to external oppression. As Reiman notes, the correlation between privacy and adulthood is already widely acknowledged (Reiman, 2004: 206-206). People without privacy, infantilised and without a vibrant inner life, won’t be easy to oppress. Rather, they won’t *need* to be oppressed, because their one-dimensional outlook will never see the need to be anything but compliant. Fried similarly argued that privacy plays a role in defending our liberty when it allows us to do or say things that are not forbidden by morality, but are nevertheless “unpopular or unconventional”. Without the privacy required to say such things on our own or in a circle of those we know and trust, we may well never do or say those things at all (Fried, 1968: 483-484). The claim is that losses of privacy will hinder the freedom of thought required to be an adult who is more than just a conformist member of a faceless, compliant herd. Without privacy, social norms would absorb the individual entirely (see also Gerstein, 1978).

Some draw a link between losses of privacy that cause conformism and the rise of autocracy. Famously, Hannah Arendt argued that a defining element of totalitarianism is the loss of privacy that attends a surveillance state (Young-Bruehl, 2008: 52-53). Edward Snowden and Julian Assange similarly link surveillance states with oppression. Snowden’s focus is government agencies such as the NSA; Assange’s concern is corporations, arguing that, “The advance of information technology epitomized by Google heralds the death of privacy for most people and shifts the world toward authoritarianism ...” (Greenberg, 2014; Assange, 2013).

A third effect of privacy invasions on autonomy is specifically upon *moral* agency. For philosophers, the link between autonomy and morality is strong. As we have seen, Kant defined autonomy as the freedom to be a moral self-legislator. Indeed, autonomy is widely regarded as “the defining characteristic of free moral agents” (Mackenzie and Stoljar, 2000: 5). What’s more, a direct link has been drawn from privacy to autonomy to moral agency. In 1976, Jeffrey Reiman argued that privacy is the means by which a society grants an individual the moral title to his own existence (Reiman, 1976). As such, privacy is an essential social practice: this is how society tells an individual that her existence is rightfully her own. Further, Reiman says that having moral title over oneself is about more than being able to control how one may act; it is also about determining which thoughts and bodily matters are to be known by others. Quite simply, if a person cannot control access to her thoughts, she will, Reiman argues, cease to regard herself as a person. For Reiman, privacy is “an especially important and good thing for human beings” (Reiman, 2004: 200). Ultimately, its value can be simply stated: “I can sum up that value [of maintaining privacy] as the protection of freedom, of moral personality, and of a rich and critical inner life” (Reiman, 2004: 209). If the internet is in some ways beginning to read users’ minds, as I have suggested, then those users, on Reiman’s argument, are beginning to lose moral title to their own consciousness.

Reiman argues that losses of privacy can take four forms. I have already described the risk of psychopolitical metamorphosis. To this he adds the risk of *extrinsic* loss of freedom, when lack of privacy makes people vulnerable to having their behaviour controlled by others (Reiman, 2004: 201). Further, Reiman identifies the risk of *intrinsic* loss of freedom, where “denial of privacy limits people’s freedom directly, independently of the ways in which it makes the susceptible to social pressure or penalties” (Reiman, 2004: 203). In this sense, privacy is not just a means to protect freedom, but is itself constitutive of freedom. Reiman gives the example of driving to destination X at time T. In the digital age this act has become more complex: “It now becomes driving to X at T and creating a record of driving to X at T.” Hence freedom has been curtailed. “I am no longer free to do the act ... *without leaving a record*” (Reiman, 2004: 204) When we think we’re being surveilled, we add the perspective of the viewer to our own perspective, leaving us with a kind of double vision. In Panopticon 2.0, the eyes

of others are internalised, whether we are in an intimate moment or on a road trip. Finally, Reiman identifies the *symbolic* risks. Privacy, Reiman argues, affirms an individual's self-ownership by granting her the ability and authority to withdraw from the scrutiny of others (Reiman, 2004: 205). In the informational Panopticon, says Reiman, people will be less likely to develop into individuals who think of themselves as owning themselves. This loss would be "incalculable" (Reiman, 2004: 206).

I agree with Reiman's arguments and examples. Without addressing his taxonomy, I am arguing that privacy is valuable for reasons of autonomy. Not in all cases, but in many cases. If our privacy is invaded, our autonomy may be compromised. Such invasions can have several effects. First, they can limit the self-creative enterprise of living our lives. Second, they can tend to encourage conformism, and hence, by the suppression of free thought, can foster totalitarianism. And third, they can stifle our moral development and identity.

So far, I have been primarily describing *individual* autonomy. To do so, I have been drawing on accounts that tend to champion privacy as an individual good (see Schoeman, 1984: 206). To an extent, these accounts suggest that privacy is all or nothing, and that a person is either in private or in public. This in turn evokes an *oikos/polis* distinction, metaphorically if not literally. On this view, privacy involves a retreat to some sort of inner citadel. As we have seen, however, privacy is highly nuanced and layered. In a digital context, a person may be in several places at once, some private, some public. She can be simultaneously withdrawn and exposed. What's more, it matters who is watching. It makes a difference whether my wife is watching, or a company, or a government agency. Here, the idea of relational autonomy becomes relevant. Standard forms of social interaction involve an interplay of people's observations and judgments. A relational account recognises that autonomy, just like privacy, only exists for an individual in relation to others. Autonomy, like privacy, is not simple, but multi-faceted. My privacy and my autonomy do not depend on my withdrawal from other people; rather, they are scaffolded by the observations and judgments of others. It is the texture and quality of the interaction that matters. It is not surveillance *per se* that constitutes an invalid encroachment upon my privacy, and which hence threatens my autonomy, but a particular type of surveillance. In

many cases, it is a *nonconsensual* surveillance, to which turn in the next chapter. As Mackenzie notes, “autonomy is not a context-invariant concept” (Mackenzie, 2014: 16).

Relational autonomy allows for, and indeed depends upon, personal connections and social bonds (Veltman and Piper, 2014: 4). In part, relational autonomy is a response to the feminist arguments that autonomy, as traditionally understood, is “coded masculine” in conceptions that are atomistic, individualistic and rationalistic. Relational conceptions of autonomy allow “that persons are socially embedded and that agents’ identities are formed within the context of social relationships and shaped by a complex of intersecting social determinants, such as race, class, gender, and ethnicity” (Mackenzie and Stoljar, 2000: 4-5). One ramification is that a privacy bound up with relational autonomy cannot be seen as the polar opposite of publicity. The two are interwoven, as befits the layered nature of an internet that enables the multiplication of place. It is worth noting also that Kant’s account of autonomy is not inconsistent with a relational conception, given that Kant argued that socialisation plays a part in fostering (or hampering) the development of autonomy (Formosa, 2013b: 202-203). What’s more, Sharon Anderson-Gold has mounted a Kantian defence of privacy, and it runs on relational lines. Traditionally, scholars have assumed that Kant’s strict prohibition on lying was attended by the equally strict necessity of truth-telling. This, it would seem, leaves no room for privacy. However, Anderson-Gold locates a Kantian argument for privacy in “the duties that we have to respect the humanity of others and to promote the moral development of future generations” (Anderson-Gold, 2010: 28). Anderson-Gold argues that reticence can be a virtue, in the form of non-disclosure about ourselves, and in the form of not responding to the faults of others (Kant, 1996a: 6:466; Anderson-Gold, 2010: 29). Hence it is the public and social aspect of privacy that is valuable: “A Kantian defence of privacy is not focused so much on individual rights or welfare as it is on the character of our public-social culture” (Anderson-Gold, 2010: 41).

The autonomy that has typically prevailed in the liberal tradition privileges the individual at the expense of the social; but I propose it is relational autonomy that matters for privacy. Relational autonomy recognises the role of autonomy for the individual *per se*, but also for the individual as *being-in-relation*. It recognises

that individuals exist only within a complex of social ties. For Mackenzie, an adequate conception of relational autonomy recognises: that humans are vulnerable and dependent rather than self-sufficient and rational; that persons are embodied and socially, historically and culturally embedded in a way that constitutes their identities; and that social conditions restricting the exercise of self-determination are unjust (Mackenzie, 2014: 21-22). To foster autonomy, these conditions must be addressed. In other words, to foster autonomy, it is not enough to address an individual's characteristics. Attention must also be paid to the wider context. Compared with an atomistic account, relational autonomy is richer and more layered, both because it is a more satisfactory account of the sort of autonomy that should be pursued as an ideal, but also because it chimes more harmoniously with privacy. After all, privacy is one of the main ways in which both a demarcation and a connection are drawn between self and society. Privacy is a means by which an individual situates herself (control) and is situated (externally-imposed access) within a society, both apart from it and as part of it.

In this way, privacy ought to be regarded as a public good. Privacy is, ironically, invaluable for democracy. As I have been arguing, privacy enables us to test out unpopular ideas, to resist conformism and to forge our moral identity, which in turn enables us to be informed and engaged citizens. As Rauhofer writes with regard to informational privacy in the internet age, "... the right to privacy should be accorded equal status as a public or community value, acknowledging that it is necessary to maintain widespread participation in and equal access to the democratic polity" (Rauhofer, 2008: 195). Autonomy and privacy, as we have seen, are intimately connected. What's more, autonomy and privacy only exist in relation to society. Hence it is relational autonomy that we ought to foster. Further, it is *relational privacy* that we ought to foster. On the internet, after all, data revealing one person simultaneously reveals many others (Fairfield and Engel, 2015). This point is clear, from the shadow profiles created by social networks such as Facebook (see chapter two), which depend upon a "privacy leak factor". Individual privacy is dependent upon an individual's community, and online social networks harness this dependency:

In an interlinked community, an individual's privacy is a complex property, where it is in constant mutual relationship with the systemic properties and behavioral patterns of the community at large ... [W]e should consider privacy as a collective concept ... (Sarigol et al., 2014: 104).

Here is another illustration of why the control model of privacy falls short: often, our privacy is determined not by ourselves, but by others. Vallor too notes that individual data is, in fact, much more than just individual data:

Information about me is *also* usually information about the others with whom I share my life, and thus to focus only on the question of whether *I* have something to hide is a profoundly solipsistic attitude to privacy concerns” (Vallor, 2016: 191).

Online, perhaps even more than offline, privacy is contingent on those around us. If one individual has her privacy invaded, the privacy of others is thereby being invaded too. At the same time, as I explore in the next section, privacy *informs* our relationships with those around us. Without the relational autonomy that relational privacy brings, ours would be a world more homogenous and more heteronomous.

IV - Relationships

Too often, privacy has been construed simply in terms of the individual. However, and individual account of privacy is inadequate, given that privacy comprises my relations with others. Privacy, by definition, is about my standing in relation to others. This is not just because I am a being-in-relation who is socially constituted, but also because privacy is a mechanism that enables me, on the one hand, to withdraw from others, and, on the other, to connect with others. The privacies I withhold matter, but the privacies I share matter too.

Above, I have been arguing that privacy matters for reasons of dignity and autonomy. A third justification is relationships. In this section, I aim to extend my arguments for a relational approach to privacy with a further claim: that without privacy, we cannot love, trust and befriend. In doing so, I follow Charles Fried, who justifies privacy on the basis of dignity and autonomy, but adds that privacy is as necessary for love, trust and friendship as oxygen is for combustion. For Fried, relationships require the voluntary relinquishment of parts of one’s inner self to another, a relinquishment that is only possible if those parts of the self are securely held in the first place, because people cannot relinquish something they do not hold securely (Fried, 1968: 480). Fried argues that our relationships depend

on the selective, deliberate, discretionary sharing of our inner selves with others. We might share nearly all of ourselves with our lovers; we might share much of ourselves with our friends; we might share a more limited part of ourselves with acquaintances. For Fried, a person who will not share herself cannot have a friendship or a love relationship. Conversely, a person who shares *all* of herself with everyone without discrimination cannot have friendships or love relationships either, because there would be no way to differentiate close relationships from distant relationships.

It is my thesis that privacy is not just one possible means among others to insure some other value, but that it is necessarily related to ends and relations of the most fundamental sort: respect, love, friendship and trust. Privacy is not merely a good technique for furthering these fundamental relations; rather without privacy they are simply inconceivable ... To make clear the necessity of privacy as a context for respect, love, friendship and trust is to bring out also why a threat to privacy seems to threaten our very integrity as persons (Fried, 1968: 477).

Unlike Fried, I have been arguing that privacy is not solely about control. It is also about externally-imposed restrictions upon access. Moreover, Fried's account of the "inner self" suggests privacy as involving a retreat from the world. I am arguing it is also about locating us in and connecting us to the world. These caveats notwithstanding, I now follow Fried to argue that privacy's value lies in part in its integral role in enabling humans to form and maintain relationships with one another.

Like Fried, Robert Gerstein gives an account of privacy that is largely individualistic, even as he too argues that privacy is necessary for relationships. Gerstein's specific argument, however, is founded on intimacy. In arguing that intimacy is impossible without privacy, Gerstein differentiates the roles of participant and observer (Gerstein, 1978). To be a participant is to be immersed in a situation and to be engulfed by it. To be an observer, however, requires distance and an objective attitude. Gerstein argues that intimate relationships are conducted by parties who are participants, and not merely observers. However, the act of participating can be altered with the knowledge that one is being observed. Becoming self-conscious, one can lose one's sense of free abandon. For Gerstein, objective judgment can corrupt a kind of ecstatic inner focus that intimacy can engender, given that external judgment involves losing the joy of being swept up in the intimacy of the personal. Gerstein's account is individual-centric, proposing that one can only have a private life by escaping the limitations imposed by social

constraints. By contrast, I am building an account of privacy that sees social “constraints” as fundamental, and sometimes, ironically, as liberating. Social ties are not what we need to escape to have privacy; rather, they help to inform and create our privacy, which is not merely determined by our choice, but also by external factors including the law, social norms and one’s own family’s traditions and practices. In my family, it may be an unspoken principle that politics is a private matter, not to be discussed in public. This norm matters, both to my condition of privacy, but also to my relationships within my family. Breaching this tacit understanding would involve a certain breach of trust. Gerstein’s arguments are persuasive. As he shows, intimacy and relationships rely on privacy. However, his account of privacy as requiring an absence of social constraints is insufficiently relational. A private life is more than just an inner citadel that represents some personal *oikos*.

James Rachels provides an account of privacy more compatible with relational autonomy, arguing that any satisfactory account must satisfy two conditions. First, it must account for privacy in normal and ordinary situations. “By this I mean situations in which there is nothing embarrassing or shameful or unpopular in what we are doing, and nothing ominous or threatening connected with its possible disclosure” (Rachels, 1975: 325). Second, it will help explain why some information is not another person’s business, and why prying is regarded as improper. To build such an account, Rachels turns to relationships. Privacy is morally valuable, he argues, because it enables persons to control the patterns of behaviour necessary for them to build stable and meaningful relationships. Rachels observes that people behave differently around different people, and what constitutes appropriate behaviour varies significantly from one relationship to another. A man behaves differently with his wife, his mother-in-law and his boss, and there is nothing dishonest or hypocritical about this (Rachels, 1975: 326-327). As Rachels notes, some have disagreed, arguing that it is phony and inauthentic for an individual to have various modes of behaviour, and that all these social “masks” hide the “real” person underneath. Facebook’s Mark Zuckerberg seems to support this view, saying in 2010: “Having two identities for yourself is an example of a lack of integrity” (quoted in Meikle and Young, 2012: 129). Not revealing information about oneself, it is argued, can be the moral equivalent of deception (see Schoeman, 1984: 211). However, Rachels argues that this line of

argument is “quite wrong”: “the different patterns of behaviour are (partly) what define the different relationships” (Rachels, 1975: 327). Moreover, in each relationship there exists an appreciation of what kind of knowledge about one another it is appropriate to share and to know. This is how Rachels satisfies the first of his two conditions: he argues that privacy is important in normal situations because it plays a key role in allowing a person to maintain a diverse range of social relationships. This then feeds into his response to the second condition. What makes a particular piece of information about you *not* the business of someone else? For Rachels, information about you is not another person’s business if nothing about the relationship entitles the other person to know this information about you.

There have, of course, been objections raised to justifications founded on relationships. Responding to Fried, Reiman writes: “I think that Fried is wrong about intimate relations, since I think that intimate relations are a function of how much people care about each other, not how much they know each other” (Reiman, 2004: 198). As Reiman notes, one can share intimate, private information with a doctor, and yet have no relationship whatsoever. True enough. On similar lines, Cocking and Kennett dismiss the “secrets view”, arguing that friendship is about trust and caring, and not about the sharing of personal information (Cocking and Kennett, 1998). But then how to explain FOMO, the Fear Of Missing Out, which has become so common that it was added to the Oxford English Dictionary in 2013 (Barker, 2016)? A 2015 survey of social media use found 51 per cent of Australian teenagers felt anxious if they did not know what their friends were doing (APS, 2015: 35). Imagine your friend has decided to get married, and has told a handful of her closest friends, but not you. Unless there are extenuating circumstances (she tried to reach you but couldn’t; she was worried about the impact of the news on you; etc), the realisation would be painful. Admittedly, the pain may well come from the sudden awareness that you value her friendship more than she values yours, and that she doesn’t care as much as you do. Hence Reiman is right to suggest that relationships are about caring. However, I propose that they are also about knowing. In part, it is through the sharing of privacies that care is expressed, and it is on the sharing of privacies that care is sometimes built. Further, it is on the sharing (and withholding) of privacies that *trust* is built. If I am in a monogamous relationship, I want my

partner to be faithful. One way to monitor such fidelity is to know everything about her by tracking every movement and communication. That, however, would be creepy. And untrusting. It also leaves no room for betrayal. It is precisely the prospect of betrayal, made possible by the fact that partners afford one another privacies and other freedoms, that makes fidelity meaningful, and thereby solidifies a relationship. Privacy is essential for strong ties, thanks both to what is shared, and what is not shared.

An important clarification is that I am not envisaging relationships as a series of concentric circles or spheres, along the lines of the onion model of privacy described in the previous chapter. If we adapt the onion model to relationships, then our closest relationships are at the centre, where we share our privacies most liberally. As we move outwards through layers, so we move through layers of friends, from close to distant, until we come to the skin, which represents our relationships with strangers, where no privacies are exchanged. This, I suggest, is too simplistic. Sometimes we keep privacies from those closest to us; sometimes we share privacies with strangers. Like Fried, Gerstein and Rachels, I am arguing that the way we share our privacies enables us to make and keep our relationships; however, I am not suggesting simply that the more we share, the closer we are. For instance, I consider myself equally close to my wife and my children, but I share very different privacies with them. Having said that, some bits of data are generally more private than others. *Prima facie*, a nude portrait is more likely to be private than a clothed portrait. However, the level of privacy of any piece of data shifts and changes according to context. For Rachels, privacy enables persons to control the patterns of behaviour necessary for them to build stable and meaningful relationships. These patterns of behaviour vary, and information is contingent on these patterns (Rachels, 1975). Hence information that is generally highly private among one's peers will not be highly private in a medical or legal context, where there is an understanding that highly intimate information is being shared dispassionately, rather than amid the charged ebb and flow of personal relationships. This helps explain lawyer/client and doctor/patient confidentiality, where the information may be intimate, but where the context is professional and the relationship is dispassionate. As Schoeman notes, the argument that the privacy of some information depends on the context of the relevant relationship is thus not incompatible with the argument that some information is inherently more

private than other information (see Schoeman, 1984: 208-209). Just as an account of relational autonomy recognises that autonomy varies according to context, so too relational privacy must recognise that privacy varies according to context.

Ruth Gavison acknowledges the individual value of privacy, but also its public value. In a conception that can be construed as relational, Gavison argues that privacy is essential for the maintenance of important relationships. At times, argues Gavison, people will disagree, and will be unable to traverse the gulf of their disagreement. They may, indeed, be intolerant of the other's values or behaviour, even while acknowledging the legitimacy of such values or behaviour. In such cases, privacy allows for interaction without the need for addressing the areas of disagreement. In other words, privacy affords *practical* tolerance in lieu of *actual* tolerance. Privacy, argues Gavison, enables people to engage productively in situations where there is profound disagreement but also a need to cooperate. In this way, privacy enables people in important relationships to maintain their individuality; and, by enabling relationships to function, privacy contributes to a harmonious society. A spouse may understand a partner's need to fantasise, even if knowing about those fantasies would be hurtful, and hence "respect for privacy is a way to force ourselves to be as tolerant as we know we should be" (Gavison, 1980: 451-452). For Gavison, privacy is both an individual and social good.

Conventionally, privacy is often defended as a private/individual good, justified on grounds including dignity and autonomy. Then there are some who value it as a public/social good, justifying it on the grounds of social links, including our relationships (eg. Fairfield and Engel, 2015). Schoeman separates these two strands, arguing that they constitute two distinct ways to justify privacy (Schoeman, 1984: 203-209). Contra Schoeman, my aim is not to polarise privacy's value as individual good here and social good there. Rather, the examples I have given show that privacy tends to be simultaneously valuable for individuals, and for society. If a woman requires privacy for childbirth, then that will benefit her, but society too. Just as a mother wants a healthy child, so does her family, her community and society. More generally, if privacy is valuable for the sake of human dignity, then it is valuable for both individuals and society. If individual dignity is not respected then individuals cannot flourish; nor can

society flourish. Simply, a society that fails to respect the dignity of its citizens is impoverished. Similarly, if privacy is valuable for the sake of relationships, then it is not just valuable as a social good, but as an individual good too. Relationships are required not just for society to thrive, but for individuals to thrive.

Online as offline, privacy is meaningful only vis-à-vis our links with others. In this sense, privacy is built upon the recognition “that persons are socially embedded and that agents’ identities are formed within the context of social relationships and shaped by a complex of intersecting social determinants” (Mackenzie and Stoljar, 2000: 4). We are fathers, sisters, daughters, lovers, friends and more. These relationships circumscribe (and also extend) our freedom, including our freedom to be private. I do not choose to be my mother’s son, and so I do not choose the obligations (and the rights) that attend being a mother’s son, and yet I am bound by (and entitled to) them. I am a social being. Hence I determine my privacy, but only within the context of my social ties, and so in part my privacy is also determined by who I am in a society, in a community, in a family. And privacy, in turn, helps to determine and forge my social ties. It enables me to make and maintain relationships. Privacy is relational, even as it is also necessary for our relationships: it is relational in that it can exist only in the context of our relationships; and it is necessary for our relationships because without it we would be unable to love, trust and befriend.

What I have aimed to show is that privacy matters a great deal, for reasons that include dignity, autonomy and relationships. However, that does not necessarily mean that privacy *always* matters. There are certainly cases where it is legally trivial (see chapter six). What’s more, when privacy does matter, I am not suggesting that it matters for all of the reasons I have cited. In one case, it may be valuable for its role in enabling autonomy; in another, it may be valuable for its role in enabling the maintenance of a friendship; in a third, it may be valuable for those reasons and more; in a fourth, privacy may be morally trivial after all. Simply, I am arguing that privacy has many important justifications, and that these justifications overlap, so that different values come to the fore in different circumstances. In this way, I am offering an account of privacy that is more relational than individualistic, that is comprehensive but context-dependent and that posits privacy as both private good and public good.

In 1980, Gavison described privacy as a fledgling right, vulnerable to legalistic and philosophical attack. This remains the case: privacy law is still *becoming*. Nonetheless, as Gavison argues, privacy represents something basic and distinctive among both social and moral values. As such, its growing legal recognition, coupled with the persistence of claims of privacy, reveal that privacy exists to protect something important, and something which other legal categories have failed to cover satisfactorily. Privacy skeptics argue that all privacy claims could be defended by appealing to other moral and legal categories, and that we would do better to eliminate all talk of privacy (Schoeman, 1984: 200). I follow Gavison to argue that privacy (like autonomy) may be tied to a complex of concepts, but is nonetheless both distinct and coherent:

The reasons for which we claim privacy in different situations are similar. They are related to the functions privacy has in our lives: the promotion of liberty, autonomy, selfhood, and human relations, and furthering the existence of a free society (Gavison, 1980: 423-424).

In 2015, UK police announced they were investigating a new crime of “cyber-flashing”, after a female commuter received two images on her phone of an unknown man’s penis (BBC, 2015). Clearly, the victim of this breach of privacy is not the person whose privacy was compromised (the man) but the person who witnessed the compromise (the woman). Here, the value of privacy lies in protecting society-at-large from one man’s desire to expose himself. Privacy norms exist not just to protect the person whose privacy is in question, but also to protect society at large. Complex and complicated as it may be, privacy, at heart, concerns how we relate to others. Paradoxically, it both separates and connects. Sometimes it does both at once. A relational account allows for such a push-pull, confirming that privacy’s value lies sometimes in its ability to isolate, and sometimes in its power to bind.

Conclusion

For millennia, philosophers (and children) have pondered the merits of privacy. I began by contemplating two imaginary worlds: one with no privacy; the other with absolute privacy. Both seem uninhabitable. Luckily, both are also impossible

(at least for now). These imaginary worlds suggest that societies and individuals have an interest in privacy, but an interest that is qualified. I then articulated more precisely why and when privacy matters. First, I showed that some invasions of privacy are violations of dignity. This can occur in cases of secret surveillance (such as webcam spying, shadow profiles and some government monitoring) and also in cases of vulnerability (such as a woman in childbirth). I then turned to autonomy, arguing that privacy violations potentially impede autonomy in at least three ways: they constrict the creative self-expression involved in living our lives; they tend to breed conformism, which has the further effect of enabling totalitarianism; and they obstruct moral development and integrity. More specifically, I tied privacy to relational autonomy, which recognises that we are all socially embedded. For a third justification, I turned to relationships, arguing that without privacy we cannot love, trust and befriend. This led me to outline a conception of relational privacy, which recognises not only that we are all socially constituted, but also that privacy is as much a public good as it is a private good. It both separates and binds. I further showed that the value of privacy is not constant, but varies according to context. Sometimes it matters for dignity, autonomy or relationships; sometimes for a combination thereof; and sometimes, perhaps, privacy is morally trivial. This context-dependency can lead to normative despair and inertia. However, in the next chapter I will argue that, as messy and complex as privacy is, we can nonetheless apply a single prescription to powerful effect. In its defence of dignity, autonomy and relationships, Kant's formula of humanity can help us to determine when and how privacy ought to be protected.

It would be a wild exaggeration to say that every small loss of privacy enslaves us. Some such losses may in fact be morally trivial. Other losses may be justified and even desirable. However, a due dose of privacy is required for individuals and societies to flourish. As an unwitting reality TV star, Truman Burbank has no privacy whatsoever. And at first glance his life seems perfect. Happily cocooned in suburbia, he navigates a quietly satisfying existence, oblivious to the fact that his neighbours are extras, his wife is paid to be affectionate and that an audience of many millions is watching his every move. Inexorably, however, Truman begins to suspect that his reality is engineered, and that his life is playing out entirely in public. His dignity, autonomy and relationships are all suffering. Seeking to escape, he is thwarted, until finally he encounters Christof, the god-

like director of the show that has been a worldwide hit for all the 30 years of his life, and of which he has been the star. “I know you better than you know yourself,” Christof says, but Truman has a comeback: “You never had a camera in my head.” This stings Christof, who, by the way, zealously guards his own privacy. For Truman, for Christof and for my five-year-old daughter, privacy matters. Even in the faux reality of a Hollywood dome, observed by hundreds of cameras, Truman had a *modicum* of privacy. Now that he’s escaped, he’s going to find himself just as much as each of us deserves.

Chapter 5

Privacy by consent

Batman is a well-known good guy. What is less well-known is that the caped crusader defeats the Joker only by breaching the privacy of millions. In the 2008 film *The Dark Knight*, the Joker has rigged two crowded ferries with explosives. To catch him, Batman and his team at Wayne Enterprises build a city-wide surveillance system that will intercept the high frequency signals emitted by the smartphones of Gotham's citizens. Purpose-built for mass surveillance, the system makes a neat metaphor for the NSA's clandestine monitoring of US citizens' mobile phones (CriticalCommons, 2016). In this way, Batman can keep watch over the whole city. "Beautiful, isn't it?" asks Batman, showing off a wall of monitors. However, Lucius Fox, his chief of research and development, is not pleased: "Beautiful. Unethical. Dangerous," says Lucius. "This is wrong." With mumbled gravitas, Batman responds that the database is encrypted and can be accessed by only one person: Lucius himself. Still, Lucius is unimpressed. "This is too much power for one person," he says. "Spying on 30 million people isn't part of my job description." Finally, Lucius relents, with a caveat. "I'll help you this one time. But consider this my resignation. As long as this machine is at Wayne Enterprises, I won't be."

The scene raises key issues about the ethics of privacy and technology. Both Batman and Lucius acknowledge it is wrong to engage in such clandestine surveillance. However, Batman is prepared to breach one ethical principle (respecting the privacy of citizens) in order to uphold another (maintaining the safety of citizens). Ultimately, Batman's view prevails, with the commitment that the technology will then be destroyed. The scene reveals three salient points. First, it shows that privacy is not just about control, but also about externally-imposed restrictions on access. Batman has invented a device akin to Judith Jarvis Thomson's ray gun (see chapter three), and so it is externally-imposed decisions about access, not the citizens' control, that will determine citizen privacy. Second, it suggests that privacy does not exist in a vacuum, but must be balanced against other rights. And third, the significance of consent, though not made explicit, is heavily implied. The key to this breach of privacy, it seems, is that Gotham's

residents are unwitting. If they knew and consented, there would seemingly be no problem. Instead, Batman has overridden their lack of consent. In the real world, this would be an act of vigilante justice, and hence ethically dubious. However, what if Batman's actions were instead authorised by the elected legislature, which was effectively expressing the collective consent of the people? Consent, I suggest, is at the heart of Batman's ethical dilemma, and it forms the subject of this chapter.

In chapters one and two, I described the way our internet use is challenging and confusing the *condition* of privacy, and is thus necessitating that we revisit the issue of the *right* to privacy. In chapters three and four, I supported an access model, before arguing that privacy is both instrumentally and non-instrumentally valuable, for reasons including dignity, autonomy and relationships. In this chapter, I argue that we can remedy many of the problems attending internet privacy by turning to Kant's formula of humanity, an ethical principle founded on dignity and autonomy. Kant's formula not only enables us to see the issue of internet privacy more clearly, but also enables us to devise effective protections. Applying the formula, I argue, involves applying consent. More specifically, it involves applying a two-tier model of consent which overlays individual consent with collective consent. First, we must ask: does user x consent to sharing her data? To do so, we need to spell out a specific conception of consent, which blends actual consent with possible consent. However, even armed with the very best conception of individual consent, we will be restricted in our progress. What we need is another, second level, which we can think of as collective consent. This is the law - as long as it is *just* law. In other words, here I make the transition from ethics to politics, to argue that this second layer of consent overarches the first layer. This second layer can work in many ways, including to buttress, mandate, qualify or override the first layer. Further, this two-tiered model of consent dovetails with the access model of privacy elaborated in chapter three, which also allows for control: individual consent is an expression of control; collective consent is an expression of externally-imposed restrictions upon access.

The chapter is divided into four sections. The first describes Kant's formula of humanity; the second discusses consent and its link to the formula; the third examines individual consent; and the fourth provides an account of collective

consent. Together, they aim to set the normative parameters for what follows. In this chapter, then, I spell out my ethical theory of the right to privacy, and how it ought to be protected, before I turn in the next, and final, chapter to applying that theory to the internet by proposing specific legal and extra-legal protections.

I – Is the formula of humanity a formula for privacy?

The formula of humanity has been described as the central normative principle of Kantian ethics (Wood, 1999: 111-155; Formosa, 2017, in press).³⁶ It prescribes, “Act in such a way that you treat humanity, whether in your own person or in the person of any other, never simply as a means, but always at the same time as an end” (Kant, 2009: 429).³⁷ Though some Kantian precepts now seem outdated, the formula’s prescriptions are still widely current in the forms: don’t use people; and treat people with respect (O’Neill, 1989: 105). It is at once elementary and radical, situating the supreme good within ourselves. It posits absolute moral worth in people’s very personhood, and not in an external entity such as god, nor in a goal such as happiness, nor in a calculation of the greatest good for the greatest number. Simply, the formula is a moral prescription based upon humanity, in promotion of the end of humanity. On the ethically-contested internet, it can shine a steady light.

In his moral theory, Kant aimed to articulate a supreme moral principle, derived from reason, that holds true for all rational beings, without exception. He called this principle the categorical imperative, which he then presented in three iterations: the formula of universal law; the formula of humanity; and the formula of autonomy. To these he then added two supplementary versions: the formula of the law of nature, which corresponds with the formula of universal law; and the formula of the realm of ends, which corresponds with the formula of autonomy (Kant, 2009: 414-437).³⁸ These formulae are commonly known by their

³⁶ The formula of humanity is also known, more accurately, as the formula of humanity as end in itself (Wood, 1999: 111) and, by Kant, as the formula of the end in itself (Kant, 2009: 427).

³⁷ For ease of reference, the page numbers I cite for the *Groundwork* are from the edition issued by the Royal Prussian Academy in Berlin, which appear in Paton’s marginalia, rather than the page numbers of Paton’s translation.

³⁸ This interpretation is standard, but not unanimous. Henry Allison argues that the FUL is the meta-law, which exists in two distinct versions, and that the FH and FA/FRE are sub-forms of the FUL (Allison, 2011: 246-260).

acronyms: FUL, FH and FA for the three primary versions; and FLN and FRE for the supplementary iterations. For Kant, the three main formulae are intended to serve different functions. Universality gives us the *form* of the moral law; rational nature or humanity as an end in itself gives us the *material* of the law; and autonomous legislation in a realm of ends represents “a complete determination of maxims” and a totality of ends (Kant, 2009: 436; see Korsgaard, 1996: 106). Intriguingly, Kant also wrote that all three of these versions are equivalent, describing them as “precisely the same law”, without justifying his claim in any detail (Kant, 2009: 436). O’Neill describes the claim as “puzzling” (O’Neill, 1989: 105). Wood, however, argues that “the supreme principle of morality is *adequately* expressed only in the *system* of all three” (Wood, 1999: 187). The equivalence issue remains controversial (Allison, 2011: 255-260; Formosa, 2017, in press). I leave it as an open question. My argument concerns the formula of humanity and its application, irrespective of whether it is equivalent to alternative formulations of the categorical imperative.

For Kant himself, the preferred approach for moral judgment involved an application of FUL: “Act only on that maxim through which you can at the same time will that it should become a universal law” (Kant, 2009: 421). However, what is the maxim, or motivating principle, that underlies an action? The answer is often unclear. In recent decades, FUL has fallen out of favour (Foot, 1972: 305; Formosa, 2017, in press: ch. 1). What’s more, even Kant acknowledged that invoking all three versions would enable us to bring the categorical imperative “nearer to intuition” and “nearer to feeling” (Kant, 2009: 436-437). Further, Kant described the FH as a “supreme practical principle” that gives us the categorical imperative “so far as the human will is concerned” (Kant, 2009: 428). For Wood, FUL and FLN yield merely the *concept* of a categorical imperative, whereas the FH is *substantive* (Wood, 1999: 97, 111). As Wood notes, a close reading of Kant’s texts, and in particular *The Metaphysics of Morals*, reveals that the FH was Kant’s own preferred formula of application: “FH is by a wide margin the formula of choice in justifying the system of duties” (Wood, 1999: 139). Wood argues that it is from the FH that Kant derives, among others, the duty against lying, the duty to develop our natural perfection, the duty to sympathise with others and the three duties of respect for others, which are the duties against self-love, contempt and giving scandal. These duties are eminently relevant to internet privacy. We can

only ever hope to understand Kant's ethics, writes Wood, by understanding how to apply the formula of humanity (Wood, 1999: 141).

There are those who argue that the categorical imperative is an exercise in "empty formalism", and who dismiss Kant's ethics as abstract, inflexible and insensitive (see Herman, 1993: 73-74; Wood, 1999: 114). I have addressed these and other charges in the introduction to this thesis. Those who are partial to Kant, however, are particularly sympathetic to the formula of humanity, which is regarded as substantial, user-friendly and appealingly intuitive (Johnson, 2004). It is "in practice a useable, coherent, and intuitively powerful principle ... We can use the FH as a moral guide to what duties and obligations we have in particular cases" (Formosa, 2017, in press). For our purposes, it is also worth noting that the FH is founded on notions of dignity, respect and autonomy. As we have seen, dignity, respect and autonomy are also among the most powerful justifications for privacy. In its defence of dignity, respect and autonomy, the FH provides an effective means for protecting privacy, including by clarifying in which cases privacy ought to be protected. It can help us to articulate the right to privacy so that we might more effectively re-establish the condition of privacy. In short, privacy and the formula of humanity are a good fit.

How is the formula linked to dignity, respect and autonomy? The most obvious link is with autonomy, given that Kant tells us that each version of the categorical imperative is equivalent, and given that Kant's explication of the formula of humanity is closely followed by his account of the formula of autonomy, which is "the Idea of the will of every rational being as a will which makes universal law" (Kant, 2009: 431). On Kant's logic, the formula of humanity recognises the absolute worth of self-governing agents whose morality is self-legislated (Kant, 2009: 435). A few pages later in the *Groundwork*, Kant comes to dignity. Kant writes that everything has a price; or, "if it is exalted above all price and so admits of no equivalent, then it has a dignity ... [and] morality, and humanity so far as it is capable of morality, is the only thing which has dignity" (Kant, 2009: 434-435). As noted in chapter four, Kant considers autonomy to be the foundation of human dignity: "*Autonomy* is therefore the ground of the dignity of human nature and of every rational nature" (Kant, 2009: 436). Hence: "An action which is compatible with the autonomy of the will is *permitted*; one which does not harmonize with it

is *forbidden*” (Kant, 2009: 439). And the type of autonomy that particularly interests Kant is *moral* autonomy. As Formosa writes, “Dignity is associated with being not merely subservient to the moral law, but being autonomous with respect to it” (Formosa, 2017, in press).

The term “respect” does not feature as prominently in the *Groundwork*, although Kant does write that the rational nature of persons marks them out as “ein Gegenstand der Achtung”, meaning “an object of respect”, or, in Paton’s translation, “an object of reverence” (Kant, 1870: 53; Kant, 2009: 428). For Wood, as for others, respect is at the heart of Kant’s ethical project, and at the heart of the formula of humanity:

Kant’s theory maintains that to act morally is always to act for the sake of a person, or more precisely, for the sake of humanity in someone’s person. In following a categorical imperative, the determining ground of the will is the objective worth of humanity or rational nature, as an object of respect (Wood, 1999: 117).

Simply, the formula of humanity commands respect for the dignity and autonomy of all persons. This becomes clear repeatedly in *The Metaphysics of Morals*: it is from FH that Kant derives our innate right to freedom, with its suggestions of autonomy; the violation of the duty of gratitude is based on “pride in the dignity of humanity in one’s own person”; and, as noted above, all three duties of respect for others (against self-love, contempt and giving scandal) are linked to the FH in that they are grounded on “the dignity in other human beings” (Kant, 1996a: 6:237, 459, 462; see Wood, 1999: 140). If our aim is to mount a defence of privacy by invoking a normative principle founded on the notion of respect for the dignity and autonomy of persons, then the FH suggests itself as a prime candidate.

For Kant, the value in humanity lies in its rational capacities, because a rational being is able to set ends for itself. As Kant writes, “Rational nature sets itself out from all other things by the fact that it sets itself an end” (Kant, 2009: 437). And elsewhere: “Rational beings ... are called *persons* because their nature already marks them out as ends in themselves - that is, as something which ought not to be used merely as a means” (Kant, 2009: 428). Humans, as rational beings, set ends for themselves; and all rational beings must treat all other rational beings as agents who set themselves ends. Here, as Wood writes, Kant gives his categorical imperative its “objective ground”, providing normative substance by prescribing

that “humanity, or ‘the human being and every rational being in general,’ is the end in itself” (Wood, 1999: 114). Humanity is the end in itself, because humanity sets itself ends. On this point, it is also worth noting that the formula of humanity allows us to account for *degrees* of wrongness. Generally, the more harm that one intends to do to one’s rational capacities, or to the rational capacities of others, the greater the wrong (Formosa, 2013a: 13-14). This is particularly significant in the context of privacy, where some violations are more grievous and systemic than others.

For Kant scholars including O’Neill (1989), Kerstein (2009) and Formosa (2017), the FH can be divided into two subsidiary principles: the mere means principle, or MMP, taken from the command to “treat humanity never merely as a means”; and the ends in themselves principle, or ETP, taken from the phrase “treat humanity always as an end”. Just how do these two normative components relate to each other? And how do they diverge?

Various responses have been suggested. For O’Neill, if I use someone merely as a means, then I fail to respect her humanity. However, there are other ways to fail to respect her humanity. This occurs if I fail to take positive measures to treat her as an end in herself, and thus to “endeavour to further the ends of others” (Kant, 2009: 430). O’Neill, it seems, is arguing that the MMP is a sub-set of the ETP, and that any breach of the former is necessarily a breach of the latter. The latter requires more of us. “Merely not to be used is not enough for being treated as a person. Making another into a tool or instrument in my project is one way of failing to treat that other as a person; but only one way” (O’Neill, 1989: 105). Imagine that, with your consent, I employ you for a menial task (say, gardening) that will keep you from fulfilling your larger goal (studying for university). Even if I obtain your consent, I may not make it my end to help you achieve your ends. In fact, I may make it my end *never* to help you achieve your ends. As O’Neill writes: “Even when others do not deceive or coerce us, or treat us in any way as tools, we may yet feel that they do not treat us as persons either” (O’Neill, 1989: 111). Contra O’Neill, Formosa argues that the two principles are distinct. Like O’Neill, Formosa argues that some breaches of the ETP are not breaches of the MMP; but he diverges to propose that some breaches of the MMP are not breaches of the ETP (Formosa, 2017, in press: ch. 3). Cases of paternalism are an

example of the latter. I know that after a big week at work my wife likes to have a quiet Friday night. However, on this one particular Friday, my wife tells me in the morning that she wants a big night out with friends. I know better, and so contrive to foil her big night. My actions satisfy the ETP by furthering her larger end (of savouring a quiet night in), but treat her merely as a means by disregarding her direct wish (to go out).³⁹ I take no stance on this debate about the precise relationship between the MMP and ETP. Rather, my position is that *both* principles need to be satisfied for compliance with the formula of humanity.

Further, I am arguing that the MMP and ETP align neatly with Kant's distinction between perfect and imperfect duties. Kant wrote that ethical duties are either perfect, in which case they command that a person do or not do a *specific* action, or imperfect, in which case they prescribe that a person adopt an obligatory *general* end. An example of a perfect duty is the prohibition on lying; an example of an imperfect duty is the obligation to promote the happiness of others. As Formosa argues, from the mere means principle are derived all and only our perfect duties, and from the ends in themselves principle are derived all and only our imperfect duties (Formosa, 2017, in press: ch. 3). For its part, the mere means principle sets out what is morally impermissible, and affords no exception. Kerstein writes, "In contemporary Kantian ethics, the mere means principle plays the role of a moral constraint: it limits what we may do, even in the service of promoting the overall good" (Kerstein, 2009: 163). By contrast, the ETP stipulates that we treat others as persons, and thereby provides a criterion for the wide duties of self-cultivation and beneficence. That is, the ETP prescribes that we be guided by the imperfect, or "wide", duties of developing one's own talents and contributing to the happiness of others (O'Neill, 1989: 105).

The ETP is difficult to enforce. As a prescription that sets imperfect duties, it is necessarily open-ended. The MMP, by contrast, is a normative principle that aims to set clear limits. It gives us perfect duties. It is, as Kerstein notes, a "moral constraint [which] limits what we may do" (Kerstein, 2009: 163). It can, as I will show, provide clarity and consistency. My focus for the remainder of this chapter

³⁹ Admittedly, this is debatable. Perhaps I am not furthering her ends, and am not treating her as an end in herself, by failing to acknowledge that her *specific* desire for a big night out overrides her *general* wish for a quiet Friday night.

will be the MMP, which raises the question: how do we know when we are treating humanity merely as a means?

II – The role of consent

I have argued above that the formula of humanity, as Kant’s “practical” iteration of the categorical imperative, is highly intuitive and eminently applicable. In this section, I propose that it is by means of consent that the formula’s mere means principle can be applied. In sections three and four, I will then proceed to articulate a two-tier model that supplements individual consent with collective consent.

The notion of consent, rather than the exclusive domain of philosophers, is a fundamental principle of law, medicine and our everyday interactions. In law, the “age of consent” is the age at which people are deemed to have the requisite emotional maturity to have sex (Dean, 2016). In medicine, the notion of “informed consent” is fundamental, leading to questions such as when implied consent can be considered sufficient (Aveyard, 2002; O’Neill, 2003; Snow and Fleming, 2014). And in common parlance, consent is a thread woven into the texture of our day-to-day. Just as the MMP and ETP are approximated by phrases such as, “Don’t be a user” and “don’t dis(respect) me”, so too consent is suggested in phrases such as, “Is she ok with that?”, or “I didn’t sign up for that”. The term itself is commonplace, including on the net, where it flourishes in “notice-and-consent” provisions. (Even if those provisions themselves are sometimes opaque and obfuscatory.) Better yet, consent is much easier to define than privacy. Clearly and simply, it can be defined as “voluntary agreement in light of relevant information” (Audi, ed., 1999: 437).

Consent, then, is a central principle of modern life. Further, many philosophers have drawn a direct link from the formula of humanity to consent. Derek Parfit writes: “We treat people as ends, Kant claims, and not merely as a means, if we deliberately treat these people only in ways to which they could rationally

consent” (Parfit, 2011: 218).⁴⁰ Similarly, Christine Korsgaard writes that one way to test the mere means principle is to ask “whether another can assent to your way of acting” (Korsgaard, 1996: 139). However, in Kant’s own work the link is not immediately apparent, given that in the formula itself Kant does not explicitly use the word “consent”, just as he uses neither “dignity” nor “respect”. Rather, consent is implied. When Kant says that I must treat others as persons, and never merely as a means, he proposes that I ought not treat people as I do rocks or walls, props or tools. Rather, I must treat people as agents capable of directing their own lives in accordance with their reason. How can I intend that other people thus direct their own lives, if my projects are going to affect them significantly? In the face of my actions, how can I act such that those others remain self-determining agents? Most obviously, by allowing them to express consent and dissent towards my projects, and then by respecting their will. If I am behaving in a way that much affects another, then I ought to enable that other, via consent or dissent, either to license or veto my behaviour:

On this view it is morally objectionable to treat others in ways to which they do not consent. To do so treats another as a thing or tool, which cannot, so does not, consent to the ways in which it is used; such action fails to treat others as persons, who can choose, and may withhold consent from actions that affect them (O’Neill, 1989: 106)

Elsewhere in the *Groundwork*, the connection between the categorical imperative and consent is made explicit. Famously, consent emerges in a section devoted to perfect duties to others, where Kant writes that we breach the mere means principle when we make a false promise to another for personal gain. When we make such a false promise, Kant writes, we:

make use of another human being *merely as a means* to an end he does not share. For the man whom I seek to use for my own purposes by such a promise cannot possibly agree to my way of behaving to him, and so cannot himself share the end of the action (Kant, 2009: 429-430).

In this passage, Kant sets out two seemingly distinct criteria by which the false promisor breaches the MMP: by precluding the possibility of consent (“... cannot possibly agree ...”); and by precluding the other from sharing one’s end. This raises a question: how do obtaining consent and sharing ends fit together? Are they distinct or the same? O’Neill, Korsgaard and Parfit argue they are two distinct criteria (O’Neill, 1989: 139-140; Korsgaard, 1996: 139-140; Parfit, 2011:

⁴⁰ Note Parfit’s use of the phrase “could rationally consent”, rather than “do consent”, in line with his choice of *possible* consent as a preferable conception.

181-182). Similarly, Kerstein argues that these two criteria sometimes yield different results, and the MMP is satisfied only if consent has been obtained *and* an end has been shared (Kerstein, 2009: 175-176). Formosa, however, argues that if one criterion is satisfied, the other must necessarily be satisfied. On Formosa's reading, either constitutes a test which reveals whether or not there has been a breach of the MMP:

... end sharing and consent so understood ask the same question: can we both will that we interact together in a certain way? And we can ask this equivalently either in terms of whether we can *share the end* of interacting together or in terms of whether we can *consent* to interact together (Formosa, 2017, in press: ch. 3).

I am partial to Formosa's argument, in part due to Kant's use of the verb "einstimmen", which translates more literally as "join in" rather than "agree". However, I will not defend this argument. Rather, I simply want to reiterate that a failure to obtain requisite consent signals a breach of the MMP, and hence the formula of humanity. In what follows, I limit my arguments to an analysis of consent, to the exclusion of end sharing, even if some would argue that there are thus some breaches of the MMP that I thereby fail to catch, let alone some breaches of the ETP.

Imagine that I say to you, "Please lend me \$100 to fix my car, and I promise to pay it back tomorrow." Now imagine that I have no intention of fixing my car, but want the money to gamble. As Kant says, due to the deception, you have not been given the opportunity to consent to my proposed action. You have been duped. Without being properly informed, you cannot possibly consent to my action. (Nor, by the by, can you share in my ends.) There has thus been a breach of the MMP, and hence the formula of humanity. Consent is thus an effective test by which to gauge whether the MMP has been breached. The formula of humanity requires us not to treat others as mere means, but to respect their reason, dignity and autonomy. To do so, we are required to respect their consent and dissent. This may well not be the only thing we must do to satisfy the formula. If we fail this test, however, we are acting unethically.

And, clearly, what we are concerned to find is *morally justifying* consent. This, I argue below, involves an application of both individual and collective consent. First, we want to locate individual consent that is *morally relevant* and *sufficient*.

We want to know which consent is “voluntary agreement” in the light of “relevant information”. Clearly, many existing notice-and-consent provisions to be found on the internet do not satisfy these criteria. Too often, agreement is more mechanical than voluntary. Too often, relevant information is hidden or obscured. Once we have articulated the strongest possible version of individual consent, we must supplement it with the strongest possible conception of collective consent, as expressed in just laws. As I sketch out this two-tier model, we will again see that privacy is not just about the individual, but about society too. For digital privacy, individual consent must be overarched by legal regulation, in a form which ought, *inter alia*, to mirror the blanket prohibitions against “misleading and deceptive conduct” found in consumer protection law (see chapter six). In this way, individual consent, flawed as it often is on the net, can operate beneath the umbrella provisos and protections of a collective consent – the united will of the people – as enacted via just laws.

III – Individual consent

I have been working with a definition of consent as the “voluntary agreement in the light of relevant information.” Unfortunately, the neat precision of this phrase is illusory. When can consent be said to be “voluntary”? What information should be considered “relevant”? In which circumstances must consent be obtained? And when does the consent obtained qualify as adequate?

The medical profession, which regards “informed consent” as a founding principle for patients undergoing procedures, has grappled with these issues. There, consent prescribes that a competent adult patient is given an adequate understanding of treatment options and their risks. A common view is that both morality and law *require* that no medical procedures be performed upon competent adults without their informed consent (Audi, ed., 1999: 437). However, ever since informed consent in a medical context first emerged in the case law in the 1950s, significant debate has arisen about the scope of the doctrine. In broad terms, there are two overarching questions. One, whose consent is required? Two, when is consent adequate? Further questions then follow. Which patients are competent to consent? How much, how detailed, and what sort of information

must be given in order that a patient's consent can qualify as informed? What sort of conditions need to prevail so that the patient understands the information? And how can it be determined that consent was voluntary, and not the result of undue influence by the physician, particularly in light of the unequal relationship of doctor and patient? (Audi, ed., 1999: 437) Issues addressed by researchers include: the value of implicit consent (Aveyard, 2002); whether silence can constitute consent (Millstein et al., 1994); and the ability of schizophrenics to consent (Moser et al., 2002). In a medical context, it has been argued, consent has significant limitations (O'Neill, 2003). Similarly, it has been argued that consent has significant limitations when it comes to internet privacy (see chapters three and six).

For now, though, individual consent remains the go-to principle, both for medicine and for internet privacy. What is required, then, is that we articulate the best possible conception of consent, in order that we can most effectively apply the formula of humanity.

i. Actual consent

Privacy, I have been arguing, is sometimes about control. When it comes to internet privacy, however, some users feel as if they have already lost control. As one study found:

Rather than feeling able to make choices, Americans believe it is futile to manage what companies can learn about them. Our study reveals that more than half do not want to lose control over their information but also believe this loss of control has already happened (Turow et al., 2015: 3).

Users want control, but feel they have lost it, and with it much of their privacy. Consent, it would seem, is a way to reinstate that control. And what could be better than *actual* consent? If I am going to treat my neighbour as a person, if I am not going to treat her as a mere means, then shouldn't I obtain her actual consent to every action of mine that affects her? If I want to chop down a tree that straddles both our properties, should I not ask first, and then respect her wishes? Similarly, if a social media site wants to share my contact list with a third party, should it not ask me if I consent?

The complication is that humans are not perfectly rational and fully informed beings with a complete grasp of every relevant fact in every situation. A neighbour may not properly comprehend that chopping down her tree will involve big expense, heavy machinery and huge inconvenience. For O'Neill, the nub of the problem is the "opacity of intentionality".

When we consent to another's proposals, we consent, even when 'fully' informed, only to some specific formulation of what the other has it in mind to do. We may remain ignorant of further, perhaps equally pertinent, accounts of what is proposed, including some to which we would not consent (O'Neill, 1989: 108).

This is particularly evident online, where data flows in ways that are unforeseen and complex (see chapter two). It is also evident in medical encounters. A doctor cannot possibly convey her full knowledge and experience in every case, if in any case. Imagine that, after examining a patient with stomach pains, a doctor suspects a likely diagnosis, but is aware that other, extremely unlikely diagnoses are possible. Will the patient's consent to a treatment plan be morally justifying only if the patient is informed of each of these possible diagnoses? The answer would seem to depend on how likely those other diagnoses are, how serious their respective implications are and what the relative treatments involve. At this point, then, the doctor must make a number of *ad hoc* decisions about what to tell the patient. Given that no doctor is omniscient, given the limitations of communication, and given that no patient is ever told everything that a doctor knows, it's fair to assume that no patient is ever "fully" informed. Does this mean that no doctor, acting amid all the complexity of the human body and its many antagonists, ever obtains proper consent? No, but it does mean that morally justifying consent ought to be sensitive to the imperfections and vulnerabilities of humans as well as the way in which our relationships, engagements and communications are imprecise, incomplete and quixotic.

For O'Neill, difficulties with actual consent can be classified into three categories, each highly relevant on the internet. The first is that it is unclear precisely what constitutes consent, and as such it is hard to determine precisely what has been consented to. The clearest instances of actual consent exist in legal and institutional contexts involving explicit formal procedures, such as signatures and oaths. Even here, however, there can be ignorance, misrepresentation, duress and pressure (O'Neill, 1989: 107). A marriage, for instance, is a legal bond, formalised

with oaths and signatures. Consent is thereby ratified. However, even a marriage can be annulled on various grounds, including if one party is already married, or was coerced, or was under the influence of alcohol or drugs during the wedding. Sometimes formal procedures can mask morally problematic instances of actual consent, as with colonial treaties signed with native peoples. In 1840, New Zealand's Treaty of Waitangi was "signed" by many (but not all) Maori chiefs, making New Zealand a British colony. Ostensibly, we have here formal consent. Closer examination, however, reveals that the Maori had an oral culture and many were, at this stage, less than fully literate. Further, various versions of the treaty exist, and many chiefs believed the oral conditions they discussed during negotiations were more binding than the printed document(s). For the English, the treaty legitimated government of the Maori; for the Maori, the signatures meant less (McKenzie, 1984: 355-365). When there are no formal procedures in place, it is even harder to determine what constitutes consent. Many nurses, for instance, administer care without seeking verbal or written consent, claiming that their patient's consent is implied (Aveyard, 2002). As O'Neill notes, consent is opaque. It is a propositional attitude that may not extend to the logical implications, the probable results, or even the "indispensable presuppositions" (O'Neill, 1989: 107). The boundaries of consent can be unclear, including on the internet.

The second category of difficulties arises, for O'Neill, with hidden coercion, or "when consent given does not match the activities it supposedly legitimates" (O'Neill, 1989: 107). Marxists, for instance, argue that workers in a capitalist society do not in fact choose to work. They may choose to work for Ford or Holden, but they cannot choose *not* to work. Similarly, women cannot choose *not* to be affected by their socially prescribed gender roles.

The outward contractual form masks an underlying coercion ... A choice between marriage partners does not show that the married life has been chosen. The outward forms of market economies and of unarranged marriages may mask how trivial the range of dissent and consent is (O'Neill, 1989: 107-108).

This, to use a Marxist phrase now found in philosophy, sociology and social psychology, is the "false consciousness" by which people delude themselves, constructing their lives on suspect reasoning, unaware of their true place in society or history (Lichtheim, 1967; Augoustinos, 1999). On this view, people believe themselves to be, and often outwardly seem to be, autonomous, even as

they are in fact constrained and exploited. Is poverty coercion? Imagine an indigent father who donates a kidney to feed his family. Is his consent compromised? In personal relationships, coercion can be particularly layered and refined. If I slam a door to win a marital argument, is that coercion? What if my slamming doors has in the past prefigured verbal abuse and violence? Coercion, it seems, must be distinguished from economic pressure and mere manipulation. In the legal and medical literature, a distinction has been drawn between coercion and constrained consent, with an acknowledgement that distinguishing the two is often difficult (Anitha and Gill, 2009; O'Neill, 2003). Indeed, there are those who ask: is genuine consent even possible within the institutions of the liberal democratic state? (Pateman, 1980: 149) I proceed on the basis that genuine consent is possible, although I acknowledge this second difficulty. Underlying coercion can be an issue. Indeed, it arguably helps explain the drive to participate in social media. When all my friends are communicating on Instagram and Snapchat, I may join in despite serious reservations about the effects on my privacy. Like the factory worker and her job, I might feel that I cannot reasonably choose *not* to be on Facebook. And this, I note below, is where a conception of relational privacy can help to recognise underlying inequities, with a view to ameliorating them.

O'Neill's third range of difficulties concerns people whose abilities to consent or dissent are limited. Such people include children and the mentally ill, or whose command of language is rudimentary. Such difficulties arise often in medical ethics, with some patients simply unable to comprehend the ramifications of potential treatments, no matter how fully and well explained. In the UK, the *Mental Capacity Act (2005)* addresses this issue by presuming that every adult has the capacity to consent, but also by protecting vulnerable people deemed unable to make their own decisions. The Act specifies who can make decisions, in which situations, and how (LegislationUK, 2017). Indeed, *many* patients are limited in their capacity to consent. As noted above, few patients have the same understanding of potential treatments as their doctor; and, in any case, communication between doctor and patient may be less than ideal. In a medical context, as in some others, "consent may be spurious even when based on average understanding and a standard ability to make decisions" (O'Neill, 1989: 108). In a sense, questions of consent are more straightforward in cases of heavy

impairment. In such cases, paternalism may not just be permitted, but required. For children, the mentally ill and others whose rational capacities are immature or impaired, we need to think in such instances about scaffolded consent, and in terms of competence to consent that is afforded via social supports (see part four, below). Again, this point suggests the requirement for relational conceptions of autonomy and privacy; and yet again, these difficulties arise often in the case of internet privacy, where children can masquerade as adults, drunks can feign sobriety and schizophrenics can mask symptoms.

Perhaps the clearest arguments against actual consent lie in deception and coercion. If I deceive or coerce you into actually consenting to my project, then ostensibly I have moral authority to proceed. If I have obtained money from you with a false promise, knowing full well I will never pay that money back, then I have, *prima facie*, obtained your actual consent.⁴¹ Deception and coercion are validated by actual consent, it would seem, even though the formula of humanity outlaws them. I return to deception and coercion below. At the least, these various examples show that people's capacity to consent is less than ideal. Sometimes people misunderstand what it is to which they are consenting. Sometimes people consent blindly, without really knowing or caring to what they are consenting. Sometimes people consent to actions that are deeply against their best interest. Actual consent would appear to be inadequate to allow for the complexity of human engagements, and to allow for the frailty, vulnerability and irrationality of human beings. As the medical literature on informed consent reveals, and as judgments on consent in sexual assault cases attest, actual consent is more problematic than it first appears (Pateman, 1980). In O'Neill's elegant phrase, actual consent is problematic due to the "opacity of intentionality". When it comes to individual consent, actual consent is a part of the story, but something is missing.

ii. Hypothetical consent

What we need is an account of consent that overcomes the opacity of intentionality. A second candidate thus emerges in the form of hypothetical

⁴¹ Or have I? Another interpretation is that I have obtained actual consent only for my false promise, whereas I have obtained no consent for what I am in fact planning.

consent. Instead of asking whether actual consent has been obtained in a given case, hypothetical consent demands that we ask whether a *fully rational* person *would* consent to a similar proposal. Theoretically, the inquiry then becomes straightforward: if a fully rational person would consent to my proposed action, then I can proceed to perform that action in the secure knowledge that I am not merely using others as means. There are clear advantages to such an approach: it removes the possibility of an actual person misunderstanding what she has consented to; it prevents persons from acting against their own best interest; and it means that people cannot consent without knowing or caring what they are consenting to (O'Neill, 1989: 109-110). An analogous approach is taken in many areas of the law, in the notion of the "reasonable person". In torts law, cases of negligence often turn on whether an individual exercised her duty of care the way the famous, figurative "man on the Clapham omnibus" would have (Miller, 1987: 171). If not, she may be found to have acted negligently. Indeed, some courts have applied a reasonableness test to consent. In 1985 the House of Lords in the UK settled an issue of informed consent by deciding that a medical patient is not entitled to be told anything his doctor chooses not to disclose, as long as a responsible body of medical professionals would sanction that choice (Miller, 1987: 171). In law, reasonableness rules. In ethics, I suggest, the standard to apply for hypothetical consent is that of a fully rational person.⁴²

This type of approach addresses the three types of difficulties with actual consent identified by O'Neill. First, what constitutes consent becomes clearer, if not entirely unproblematic, precisely because consent now operates on a hypothetical level (O'Neill, 1989: 110). In a medical context, we ask not whether a patient does consent, but whether a fully rational patient would consent, thereby dispensing with all the potential problems of miscommunication, misunderstanding, misinformation, and so on. Second, consent now matches the activities it legitimates, thereby overcoming the challenge of false consciousness. Yes, the factory worker has given actual consent to be so employed, but her consent, seen in its larger social context, may be revealed to be irrational (and perhaps even

⁴² Clearly, the *reasonable* person of the law is far removed from the *fully rational* person I am proposing for hypothetical consent. The reason I prefer the latter is that: it demands more of the person seeking consent; it is less vague; and it is, as we shall see, the model preferred by philosophers such as Michael Smith (1994). Of course, a reasonableness test for consent is also possible. In any case, I will go on to argue against hypothetical consent, except in cases of incompetence.

unreasonable). In such a case, her actual consent can be overridden by her hypothetical dissent to show that she is indeed being used merely as a means. Third, hypothetical consent overcomes problems for persons with impaired capacities to consent. If my capacity to consent is limited, then a fully rational person can be posited in my place to consent on my behalf. Instead of actual consent, the test becomes consent by proxy, where the proxy is fully rational.

What then is the problem? There are at least three. The first problem is that none of us is a fully rational being, and none of us is infallibly reasonable. This matters because the irrational, unreasonable within us is, in a significant way, what defines us as human beings. If we invoke the notion of hypothetical consent, then, in a sense, we are talking about consent divorced from real life. It becomes theoretical and largely meaningless. The second, related problem is that it is the *particularity* of our individual rationality that is apposite. If hypothetical consent is the test, then each person's own specific humanity becomes insignificant. To ask, "Did she consent?" is problematic, as we have seen, but at least it allows for a specific "she". To ask, "Would she have consented, were she fully rational?" is to override that person's self-determination with a sort of moral tyranny. Via hypothetical consent, we might override actual dissent, coercing a person against their will. If I wish to borrow my neighbour's car, but she dissents, then the actual consent model stipulates that I need to respect that dissent to satisfy the MMP. However, she may have irrational reasons for dissenting. Perhaps a horoscope warned her not to trust friends. Were she fully rational, I know she would consent. Can I justifiably override her actual consent? It seems not. Whereas actual consent is an expression of her will, hypothetical consent overrides her will. As O'Neill writes, "It seems implausible that treating others as persons should even sometimes be a matter of overriding what others as we know them actually choose" (O'Neill, 1989: 109). Finally, a third problem concerns the difficulties of determining the standards that comprise "fully rational". Can these standards incorporate the consenter's desires at all? Do they take account of local values? Or must they be universal? If actual consent is troubled by the opacity of intentionality, then hypothetical consent is troubled by its idealised abandonment of irrationality, by its moral tyranny, and by the difficulties involved in defining and applying the elusive phrase "fully rational".

One version of hypothetical consent is proposed by Michael Smith. On Smith's view, morality is informed by desires. That is, Smith builds a Humean account, departing from Kant to argue that actions are not motivated by reason alone. Hence, given that our desires can be self-destructive or anti-social, our moral decisions are fallible (Smith, 1987). As an alternative, Smith then constructs an idealised possible world with conditions of full rationality. In this world, our desires are ideal and fully coherent. Smith calls this the "evaluating" world, and proposes that it is from the perspective of this world that our actions can be evaluated. In this way, correct moral pronouncements can be made, given that a person in this world is "beyond reproach, from the point of view of reasoned criticism" (Smith, 1994: 173). Smith then constructs the notion of the fully rational adviser, who provides an answer to the question: what would I advise myself, if I were fully rational? Hence Smith proposes a fully rational version of each specific person, rather than a fully rational everyman. "In their own worlds fully rational agents will find themselves in quite different circumstances from each other, circumstances that are conditioned by their different embodiments, talents, environments and attachments in their respective worlds" (Smith, 1994: 173). My idealised rational adviser is not the same as yours, and does not necessarily proffer the same advice. Smith's account goes some way towards responding to the three problems outlined above: by allowing for individuality, it allows for a personalised rationality, thereby mitigating the prospect of moral tyranny. The prospect of tyranny still exists, except that I am being tyrannised by a fully rational version of myself, rather than an idealised, detached agent of pure reason. This may well be something of an improvement, but the phrase "fully rational" remains problematic. While Smith's rational adviser is philosophically a neat solution, and while ethically it might provide guidance, for policy-makers and legislators it would likely prove to be too fluid a notion to form the basis of any regulatory protections for privacy.

There is a further objection to Smith's rational adviser. Indeed, it is an objection that might apply to actual consent as it does to Smith's account. And that is: both seem to be intertwined with desires. Actual consent often hinges on desires, given that people's choices are usually in line with their wants or preferences; Smith's view describes actual preferences onto which rational ordering is hypothetically superimposed. On the interpretation of the formula of humanity I have outlined,

the mere means principle and the ends in themselves principle are distinct. However, there is significant overlap. In most cases, it's fair to say, if I fail to respect your consent and thus treat you as a mere means, I am also failing to treat you as a person. However, if our conception of consent essentially seeks to avoid what people don't want (the actual consent model) or to avoid what they would not rationally want (the hypothetical consent model), then it's unlikely that treating others as persons can be of prime moral importance. As O'Neill writes:

In a moral theory in which wants are basic, the notion of treating others as persons carries no independent weight. In Kantian terms we might say that the notion of a person does not matter in a heteronomous moral theory. If wants or rationalized preferences are morally fundamental, consent is of derivative concern. It is only within moral theories for beings who can sometimes act independently of desires or preferences that the notion of consent carries independent weight. In such theories it is important that consent be possible for others, but of less concern whether what they consent to is what they want (O'Neill, 1989: 112).

Kant's vision of autonomy and morality presupposes that agents can choose to act *against* their desires, including the desires that their rational selves would have. In his famous example of the sympathetic man, Kant proposes that it is the man who performs a good deed when he would rather not whose action can be clearly recognised as having moral worth. Kant writes: "The worth of character begins to show [when] he does good, not from inclination, but from duty" (Kant, 2009: 398-399; see also O'Neill, 1989: 117). To apply the formula of humanity, we require a type of consent that can be decoupled from desire. Below, I will argue for an account of actual consent that *can* be decoupled from desire. However, consent and desire cannot be prised apart in Smith's account of hypothetical consent. Unsurprisingly, Smith's Humean solution seems unsuited for our Kantian project. This leads us to possible consent.

iii. Possible consent

If actual consent suffers from the opacity of intentionality, if hypothetical consent suffers from moral tyranny and is bound up in desire, how can we do better? The solution offered by Kant, and by O'Neill, Korsgaard and Formosa, is possible consent.⁴³ This conception provides that, if our goal is to treat others not merely as a means (which usually overlaps with our goal to treat them as persons), then

⁴³ Which is not to say that they all support the same conception of possible consent. Formosa's account of possible consent, for instance, is much nearer actual consent than the others' (Formosa, 2017, in press: ch. 3).

we are required to give those others the *possibility* of consent or dissent whenever our proposed actions significantly affect them. As O'Neill writes: "The morally significant aspect of treating others as persons may lie in making their consent or dissent *possible*, rather than in what they actually consent to or would hypothetically consent to if fully rational" (O'Neill, 1989: 110-111). Simply, if the other *cannot* consent to what I propose, then I treat her merely as a means. "An agent treats another merely as a means and thus wrongly if in his treatment of the other the agent does something to which the other cannot consent" (O'Neill, 1989: 113). Similarly, Korsgaard writes, "The question whether another can assent to your way of acting can serve as a criterion of judging whether you are treating her as a mere means" (Korsgaard, 1996: 139). There is thus a profound change in focus. From the issue of *does* the other assent in the actual consent model, or *would* the other assent in the hypothetical consent model, we have shifted to the issue of *can* the other assent.

Kant himself uses the language of possibility in his example of the false promisor. When I make a false promise to another for financial gain, writes Kant, I "make use of another human being merely as a means to an end ... For, he whom I want to use for my purpose by such a promise *cannot possibly agree* to my way of behaving toward him ..." (Kant, 2009: 429-430; italics mine). The choice of adverb is subtle but significant.⁴⁴ Due to my deceit, the lender cannot *possibly* agree (or consent) to my action. The lender does not know that I want her to be a donor, not a lender. The adverb reveals Kant's preference for possible, not actual, consent. How, though, can we come to a working conception? O'Neill makes two points. One, morally significant consent is not consent to every aspect of another's proposals that may affect me. And two, we must map out what it is that makes genuine consent possible (O'Neill, 1989: 110-111).⁴⁵

O'Neill's first point is that morally significant consent cannot be consent to each and every aspect of another's proposals that may affect the potential consenter. We must endeavour to separate significant, morally relevant consent from spurious, morally trivial consent. Hence, to avoid using someone as a mere

⁴⁴ In fact, Kant writes "kann unmöglich", which translates as "can impossibly" (Kant, 1870: 54). However, the translation "cannot possibly" is both faithful and more elegant.

⁴⁵ O'Neill also makes a third point: if another's consent is to be morally significant, it must be genuine, not spurious. It must indeed be his or her consent. That is, the consenting cannot simply be up to the initiator of the action. I address this point below, in part v.

means, we need to identify the “morally significant aspects of plans, proposals and intentions” (O'Neill, 1989: 109). To borrow an example given by Kerstein: if a jogger in a park hears a stranger singing and enjoys the melody, does the jogger need the stranger's consent to keep listening? (Kerstein, 2009: 165) Conversely, does the singer need the consent of the jogger to keep singing? The jogger and the stranger are affecting one another, but seemingly neither has power of veto over the other's actions. It seems that we can only have dissent's power of veto over *some* of the actions that affect us. Formosa writes that if I intend to punch you, or lie to you, or steal from you, then the requirement for consent gives a power of veto. The situation is more nebulous, however, with noise levels. If I wish to play my radio at a volume only just audible outside my walls, does my neighbour have power of veto? Seemingly no. What about if I am planning to have an all-night jam with Marshall stacks turned to 11? Seemingly yes. As Formosa writes,

Intuitively it seems that in such cases the other's consent is needed only if the noise-making is not, at least depending on the context, merely annoying, but far more intrusive or even, at the extreme, constitutes a significant harm (a sonic assault on you akin to a punch to the ears). (Formosa, 2017, in press: ch. 3)

In cases of noise-making, there appears to be some inherently vague point, somewhere between mere annoyance and sonic battery, where the requirement to obtain others' possible consent kicks in. Is this degree of vagueness a fatal flaw in the application of the formula of humanity? No, argues Formosa, and I agree. It demands judgment and the application of discretion. What's more, as Formosa notes, a solution exists in the law. Such law could, for instance, specify that people cannot exceed a prescribed decibel level, or must keep volume “at reasonable levels”, thereby effectively resolving issues of vagueness.

The task of identifying the “morally significant” aspect of plans is especially pressing in a digital context. If I take a photo of my children outside the Sydney Opera House, it would be absurd for me to seek the consent of each one of the 100 or so people in the background, even though I may upload the photo to the web, where my photo becomes “greased data” and could possibly be seen by millions. Now imagine that when I look closely at the photo later that evening I notice that in the background a young couple is kissing passionately. Further, if I crop my photo judiciously, I have a modern version of Robert Doisneau's famous photo, *Le baiser de l'hôtel de ville* (*Kiss by the Hotel de Ville*). Unwittingly, I have

snapped a masterpiece, *Kiss at the Sydney Opera House*, which I am eager to exhibit, publish and sell. But what of the amorous couple, who happen to be clearly recognisable. Do I need their consent? Is their consent morally significant? The answer would seem to be yes, given the potential effect of my action upon the couple. Imagine that they are having an extramarital affair. Or that they are not, but the image suggests they are. Or simply that they want to keep their affection secret.

The power of Doisneau's image lies largely in the frisson of danger and desire arising from the tension between public and private. In central Paris, the town hall square is one of the city's most famous, most public spaces; the kiss, however, is intensely intimate and personal. In a way, in its confusion of public and private, the image prefigures the way the internet enables a multiplication of place, as described in chapter one. The story behind the 1950 photo is illustrative. Until 1992, it was thought that the kissing couple had been caught in an unguarded moment of public intimacy. Doisneau, renowned as a chronicler of street life, encouraged this belief by keeping the couple's identity to himself for four decades. Only in response to a court action did he reveal that the image was posed. First, Doisneau had seen two strangers kissing; then he had asked them to repeat their kiss for his lens. As he said, "I would never have dared to photograph people like that. Lovers kissing in the street, those couples are rarely legitimate" (BBC, 2005). Having agreed to participate, the couple recreated the scene at three Paris locations, whereupon an image taken at the hotel de ville appeared in *Life* magazine, before becoming famous as a poster in the 1980s. Had Doisneau shot without asking, he would have been breaching a French prohibition on taking a person's picture without their consent. As elsewhere on the Continent, French law has long held that persons in public may be photographed, but that no photograph may be published that focuses on them as individuals, unless they consent (Whitman, 2004: 1197). Indeed, it was under such a law – or by this enactment of "collective consent" – that Doisneau was sued by a pair falsely claiming to be the lovers pictured. Only then did he reveal the couple's true identity (Henley, 2005).

By contrast, Doisneau did not obtain the consent of the passers-by in the background, at least one of whom is clearly recognisable. The salient distinction seems to be that the couple was engaged in an act often considered private,

whereas the passers-by were merely ambling in public. The former were engaged in a stolen, secret moment (or a recreation of a stolen, secret moment); the latter were behaving as anyone might behave in public. Doisneau was right to obtain the couple's consent. Similarly, for *Kiss at the Sydney Opera House*, it might not be easy to identify and locate the couple, but it seems that I ought to. I ought to give them the possibility to consent or dissent to my project.⁴⁶ By contrast, I would not need to do the same for any others in the background. In O'Neill's formulation, consent must pertain to the *morally salient aspect* of another's proposals; it must attach "to the deeper or more fundamental aspects of another's proposals" (O'Neill, 1989: 110). The connection between the act and the person acted upon cannot be tangential and trivial, but must be direct and significant. The distinction between the two depends on our judgment, which can be informed by rules of moral salience (see chapter six), and which can certainly be guided by the collective consent of the law. If I want to publish *Kiss at the Sydney Opera House*, the couple's amorous embrace may be seen by a wide audience. Their consent is required if I want to exhibit, publish or sell the photo. By contrast, their consent is not required as to whether I shoot them in portrait or landscape format.

O'Neill's second point is that we must map out what it is that makes genuine consent possible (O'Neill, 1989: 111). What are the required conditions? Indeed, can one set of requirements unfailingly combine to give us genuine consent? For Christine Korsgaard, possible assent requires that the potential consentor has a knowledge of and some power over the relevant events: "knowledge of what is going on and some power over the proceedings are the conditions of possible assent; without these, the concept of assent does not apply" (Korsgaard, 1996: 139). These criteria are logical, but what knowledge, specifically? How much power? O'Neill acknowledges the difficulty of identifying one set of requirements that apply in every case, but does venture that there may be *some* necessary conditions without which genuine consent or dissent is impossible. "If we coerce or deceive others, their dissent, and so their genuine consent, is in principle ruled out. Here we do indeed use others, treating them as mere props or tools in our own projects" (O'Neill, 1989: 111). In the face of deception or coercion, it seems,

⁴⁶ Or, at the least, I ought to make a reasonable effort to identify and locate them and give them the possibility to consent. But what if I cannot identify or locate them? And when does silence amount to consent? These are testing issues that are beyond the scope of this thesis, other than to say they are prime candidates for the collective consent of legal regulation.

genuine consent is precluded. As O'Neill writes elsewhere, perhaps the best reason for adopting the principle of informed consent in medicine and beyond is that it takes seriously the prospect that people are being neither deceived nor coerced (O'Neill, 2003).

As we seek to spell out the conditions that make genuine consent possible, this point is worth investigating. The mere means principle seems to suggest that coercion and deception are always wrong, given that people cannot possibly consent to being deceived or coerced. Indeed, Korsgaard writes that, "According to the formula of humanity, coercion and deception are the most fundamental forms of wrong-doing to others." However, as Parfit notes, there are patent exceptions (Parfit, 2011: 178-179). If someone is unconscious, they cannot consent to life-saving surgery, but presumably such surgery is not wrong, even though coerced. People might also freely consent to being later coerced, as some once did ahead of painful surgery in pre-anaesthetic times. Moreover, people freely consent to being legally coerced. We (mostly) agree, under threat of punishment, to pay taxes, obey speed limits and refrain from indiscriminate killing. In certain circumstances, coercion will be permissible, and even required. These cases would suggest that not all coercion is wrong, whether that coercion be from other individuals, or the state.

Perhaps deception isn't always wrong either. In certain cases, I might tell you a lie to save your life. This deception seems to be permissible, if not required: "My life-saving lie would be like life-saving surgery on some unconscious person" (Parfit, 2011: 178-179). Kant addressed this issue with his example of the killer at the door (Kant, 1996b: 8: 425-427). You are at home with a friend who has taken refuge from a killer when you hear a knock. Opening the door, you find the killer, who asks if your friend is inside. Answer honestly, and the killer will presumably murder your friend. Lie, and your friend will be spared. Is lying permissible here? Kant says no, because we have an unconditional duty (of right) not to lie. However, several Kantians disagree with Kant's argument, including Formosa:

This seems like a 'clean hands' policy gone mad. Do your duty, tell the truth (and if you don't you will be held legally accountable), and morally wipe your hands of the outcome (even if that outcome is the murder of a friend to whom you have offered refuge). Such a view seems morally repugnant (Formosa, 2008: 162).

Elsewhere, notes Formosa, Kant argues that lying is sometimes permissible. In *The Metaphysics of Morals*, Kant writes that we have merely a juridical duty not to lie, and only when lying “directly infringes upon another’s right” (Kant, 1996a: 6: 238-239; Formosa, 2008). Having concluded that it is *juridically* permissible to lie, Formosa then asks whether it is *morally* permissible too, and again finds that it is: “... respect for humanity sometimes *requires* that we lie” (Formosa, 2008: 164-165). If a lie expresses respect for the value of human dignity, as lying to avert a murder does, then such a lie is required. After all, it is dignity, not truthfulness, that has absolute worth (Formosa, 2008: 166). Korsgaard reaches the same conclusion, via different reasoning, arguing that the formula of humanity is merely an ideal. Hence, she writes: “where the attempt to live up to it would make you a tool of evil, you should not do so” (Korsgaard, 1996: 153). I return to the case of the killer at the door below, in the discussion of competence.

Let us accept, then, as Parfit, Formosa, Korsgaard and perhaps even Kant do, that sometimes coercion and deception are permissible. How can we reconcile this position with our requirement for possible consent? In our attempt to elaborate the requirements of genuine consent, how can we allow that sometimes we are not using someone merely as a means, even as we seemingly fail to offer them the possibility of consenting to our proposed action? How can we defend performing CPR on an unconscious person? How can we justify lying to a killer at the door? One response is to turn to collective consent. By calling on the law, we can establish a sort of consent by proxy in such cases, where there is no possibility of consent. As I argue below, collective consent must indeed supplement individual consent. This, however, is not the whole solution. To fill in our account of individual consent, we must also address the issue of competence, before finally spelling out a conception of individual consent that *blends* actual and possible consent.

iv. Competence

Each conception of consent explored thus far, it seems, is somehow flawed. Actual consent, beset by the opacity of intentionality, is inadequate to allow for the complexity of human engagements, and inadequate to allow for the frailty, vulnerability and irrationality of human beings. Hypothetical consent, a kind of

moral tyranny, overrides the autonomy of the individual, and her ability to set her own moral course. And possible consent, it seems, fails to allow for permissible cases of coercion and deception.⁴⁷ In this part, I build on my analysis to argue that the application of consent must begin with the issue of competence. Indeed, competence allows us to overcome the obstacles that we encounter if we rely exclusively on actual, hypothetical or possible consent. It does so by allowing us instead to rely on a blend of all three. In practice, I propose that competence be employed as a threshold issue, which then determines which type of consent is required. If a person is incompetent to consent, then hypothetical consent is required. If a person is competent to consent, as I discuss in the next part, then the best conception of consent is actual consent, so defined as *also* to incorporate the best elements of possible consent.

This approach aligns with Formosa when he argues that *possible* consent, when required, necessitates actual consent (or the absence of actual dissent) except under three conditions. (1) When someone cannot give actual consent, such as when they are unconscious. Formosa calls this “the *competency* condition”. (2) When actual dissent does not revoke an existent authorisation, as when a policeman arrests someone for a crime despite their dissent. This is “the *rationaly required* condition”. (3) When actual consent does not grant an authorisation, such as consenting to be used as a slave. This is the “*rationaly forbidden* condition” (Formosa, 2017, in press: ch. 3). In my scheme, I allow for conditions (2) and (3) in many cases by turning to the notion of collective consent of the law (which I address in the next section). It is the law that empowers a policeman to make an arrest, just as it is the law which outlaws slavery.⁴⁸ However, in the absence of just laws, it is the formula of humanity that determines what is rationally required or forbidden. In the remainder of this section, I will address the question of competence, before spelling out a prescription for an interplay of actual, possible and hypothetical consent. In the next part, I will then argue for a conception of actual consent that incorporates possible consent. My approach follows Formosa in substance but is nominally different: where Formosa

⁴⁷ Formosa’s conception of possible consent does allow for the fact that coercion and deception are sometimes permissible (Formosa, 2017, in press: ch. 3). I return to this point below.

⁴⁸ Of course, slavery also breaches the moral law. Any law that bans slavery is merely affirming and enforcing the moral law.

advocates possible consent as necessitating actual consent, I advocate actual consent, defined also to incorporate possible consent.

A respect for *competence* to consent follows naturally from the formula of humanity. As we have already seen, the wellspring of the categorical imperative is human reason. Logically, then, *competence* to reason is morally significant. If, for instance, against her will I am compromising someone in her ability to reason, and thus making her incompetent to consent, I am breaching the formula of humanity. I am treating someone as a means and not as an end if I give her a stupefying drug, or concuss her, or deliberately and dramatically confuse her so that she is unable to think straight.⁴⁹ Kant argues that we must not render useless our own rational powers. While Kant praises the moderate consumption of alcohol as a social lubricant, for instance, he also argues that we must not drink to the point of stupefaction, because by doing so we temporarily destroy the reason which is the source of our dignity (Denis, 2012: 99-100). In short, if reason is the grounding principle of the formula of humanity, and consent is an effective test to see if the formula has been breached, then it follows that competence is required for consent to be morally justifying. This leads Formosa to his competency condition, which stipulates that possible consent requires actual consent, except in cases of incompetence. This he then combines with the rationally required and rationally forbidden conditions, as outlined above (Formosa, 2017, in press: ch. 3).

Like the notion of consent, the notion of competence is widely understood, if not uncomplicated. It denotes, in broad terms, an ability to manage oneself and one's affairs. *The New Shorter Oxford English Dictionary* defines competence as: "Power, ability, capacity, (*to do, for* a task, etc.); *spec.* legal authority, qualification, or admissibility, right to take cognizance" (TNSOED, 1993: 459). In medicine, competence is a key component in patient consent (Snow and Fleming, 2014: 486). A vast medical literature covers the competence to consent of, *inter alia*, the elderly, schizophrenics and HIV patients (Moser et al., 2002; Stanley et al., 1984; Grisso and Appelbaum, 1998). In various jurisdictions, laws define competence in medical contexts, with prescriptions including that the

⁴⁹ What if, however, she has consented for me to stupefy her? What if, for instance, she has encouraged me to administer a reason-suspending drug? Am I justified in so doing? This is a difficult issue, which I will not explore. Rather, I am arguing simply that the formula demands a respect for reason, and hence a respect for the competence to reason.

patient: must understand, retain and believe information about treatment options; is able to weigh information to make a decision; and can communicate that decision. Nonetheless, the assessment of competence is often complex (Snow and Fleming, 2014: 487-488). In the law, meanwhile, competence to consent is widely presumed. A common law presumption of competence to give, or refuse, consent is commonly traced to an English judgment delivered in 1992 (Skegg, 2011: 165). There are suggestions, however, that competence operates differently in different domains. In the law, it has been argued, a presumption of competence to give or refuse consent is entirely appropriate; in medicine, however, such a presumption is arguably less appropriate. In an emergency room, for instance, it would be dangerous to presume a survivor of a suicide attempt is competent, even if she is lucid and coherent (Skegg, 2011: 187).

In their oft-cited text *Principles of biomedical ethics*, Beauchamp and Childress argue that competence is one of the five elements of informed consent commonly identified in the legal, regulatory, philosophical, medical and psychological literature. The others are disclosure, understanding, voluntariness and (in a vexatious circularity) consent itself (Beauchamp and Childress, 2001: 79). With these building blocks, they then propose a definition of informed consent:

One gives an informed consent to an intervention if (and perhaps only if) one is competent to act, receives a thorough disclosure, comprehends the disclosure, acts voluntarily, and consents to the intervention (Beauchamp and Childress, 2001: 79).

Despite disagreement as to specifics, competence is widely regarded as significant, and sometimes decisive, in matters of consent.

Previously, after describing the flaws of actual and hypothetical consent, I noted that possible consent appears to be inadequate when deception and coercion are seemingly warranted. One such case arises when someone is unconscious and requires life-saving medical treatment. Imagine, as Kerstein does, a jogger who has passed out and requires cardio-pulmonary resuscitation (Kerstein, 2009: 174). With no one around, I administer CPR and save his life. However, there is no actual consent, and I have not offered the possibility of consent. On O'Neill's formulation, he had no chance to modify or avert my course of action; on Korngard's formulation, he had absolutely no knowledge or power over my action. *Prima facie*, my failure to give the possibility of consent means that I have

used him merely as a means. This seems wrong. As Kerstein writes, “it seems wildly implausible to contend that your attempt to save him was morally impermissible” (Kerstein, 2009: 174). Kerstein’s response is to develop a Reinforced Hybrid Account in which the MMP is applied via a two-part test (involving a consent test and also an ends-sharing test), which is itself then qualified if the other is in turn breaching the MMP. This account is, in Kerstein’s words, “somewhat intricate” (Kerstein, 2009: 176). A second approach, advanced by Korsgaard, is to argue that sometimes we can excuse ourselves from strict adherence to the formula of humanity in certain cases, such as when adherence would make us a tool of evil (Korsgaard, 1996: 153). These approaches seem counter-intuitive and/or complicated. A better approach, I suggest, involves the recognition that the jogger is incompetent to consent, and hence consent for life-saving treatment must be given on his behalf, in this case by me.⁵⁰

Similarly, consider a person who has overdosed on opioids. When the ambulance arrives, paramedics administer a life-saving injection to reverse the drug’s effects and restart breathing. However, once the patient revives, she is upset and abusive. The intervention has saved her life but killed her high. Clearly, the overdose victim has never been given the possibility of consent. What’s more, judging by her reaction (which, anecdotally, is not uncommon), it seems doubtful that she would have given actual consent. Nonetheless, it appears that she has not been treated merely as a means. Indeed, Kant explicitly prescribes that we have a perfect duty not to kill ourselves (Kant, 2009: 421-422, 429). As noted earlier, perfect duties align with the MMP, meaning that committing suicide involves treating oneself as a mere means. Intervening to prevent a suicide is perhaps not just permissible, but required. Is the ethically right response different if the overdose was deliberate or accidental? It would seem not.⁵¹ More to the point for our purposes is that similar scenarios might involve intrusions upon privacy. If I visit a friend only to find her collapsed, I would be justified in checking her medicine cabinet if I suspected an overdose, in order to save her life. In such cases, we are treating people not merely as means if we prevent them from dying. And yet in all these cases consent is impossible.

⁵⁰ Following Formosa (2017), we can also argue that performing life-saving CPR is rationally required by the moral law.

⁵¹ I will not address the very complex and very important issue of euthanasia.

Given that these life-saving actions are performed on the assumption that people want to continue their lives, or at least on the Kantian principle that they are rationally and morally required to continue their lives, how can we reconcile such cases with consent? There are two answers, depending on circumstances. Either we wait until they become competent, as when someone is drunk. We wait, I suggest, only if it is reasonable to do so. If a person is drunk and we propose to encroach upon their privacy, or borrow their car, or have sex, we must wait. If the person is in mortal danger, we cannot wait.⁵² Alternatively, if a person is incompetent and it is unreasonable to wait, then hypothetical consent is required. A parent can consent on behalf of a child and a carer can consent on behalf of a mentally ill person, by asking: what would this child/person do if fully rational? Clearly, in such cases actual consent or possible consent is out of the question. In our attempt not to use these people merely as means, consent must be imputed in their stead.

Others have taken another approach, arguing that there are values other than consent at play in life and death situations. These values include trust. In medicine, where procedures can be hard to understand for patients, trust is especially important (O'Neill, 2002: 141-164). As O'Neill writes: "Autonomy has been a leading idea in philosophical writing on bioethics; trust has been marginal. This strikes me as surprising ... Trust is surely more important, and particularly so for any ethically adequate practice of medicine" (O'Neill, 2002: ix). Clearly, there is merit to such an approach. Consent, no matter how well conceived, can never hope to be the sole ethical test in all situations where one person's actions significantly affect another. Other values are required, as I discuss in chapter six, and this seems particularly true when lives are at stake. For one thing, consent is tied to the mere means principle, and the formula of humanity also prescribes the ends in themselves principle, which tells us always to treat others and ourselves as persons. It mandates respect (Parfit, 2011: 211).

Above, I recounted the case of the murderer at the door, where it seems that lying is permitted, if not required, by the formula of humanity. But why, specifically? One approach is to invoke the ETP by arguing that I am treating my friend as an

⁵² I have merely outlined the issue of who qualifies as competent to consent. Similarly, I merely sketch out the two options of waiting for someone to consent, or supplying hypothetical consent on their behalf. For a more detailed account, see Formosa (2017, in press: ch. 3).

end in herself if I lie. This is problematic, however, since perfect duties (such as the duty not to lie) trump imperfect duties (such as the imperfect duties that flow from treating my friend as an end in herself), rather than the other way around. What's more, it seems that imperfect duties cannot *require* a person to do something specific (although this has been debated) (Formosa, 2017, in press). A better approach holds that the killer has laid down violent, coercive norms in his behaviour, which entitle me to employ similar terms of interaction, including lying, in order to protect rights. It's not that the killer's consent is not required. Rather, by behaving violently, the killer has, in his mode of acting, already consented to an interaction based on coercive terms (Formosa, 2017, in press).⁵³

Further, I am sympathetic to arguments that such a case can also be seen as concerning the trust between my friend and me, particularly given that I am proposing an account of autonomy and privacy that is more relational than individualistic. My claim here, however, is more basic: when consent is at issue, first comes the question of competence, then the question of consent. For the killer at the door, consent is relevant, but the killer's consent can be disregarded. We have not given the killer the possibility of consent, and yet we have not used him as a mere means, even though it might seem at first glance that we have. In the case of the unconscious jogger, by contrast, competence is the threshold issue. Clearly the jogger is not competent to consent. Hence hypothetical consent becomes the test (unless it is reasonable to wait until he is competent, which is clearly not the case here). Meanwhile, the collective consent of the law may well have a significant supplementary role to play. The law can, for instance, compel individuals to render assistance in certain circumstances.

Prima facie, by precluding the possibility of consent, the case of the unconscious jogger and the case of the killer at the door are acts of coercion and deception that violate the mere means principle. Logically, though, the notion that we are using

⁵³ There is a second approach that authorises lying. This approach is to argue that there is no perfect duty not to lie in this case. As Wood writes, an act is required by a perfect duty if the failure to perform it would amount to a failure to respect humanity as an end in someone's person (Wood, 1999: 325). In this case, lying is required because it expresses respect for the dignity of my friend (Formosa, 2008: 166). At first glance I am treating the killer as a mere means; on closer inspection, I am avoiding treating *humanity*, and specifically the humanity in my friend, as a mere means and, in this case, this is the best way to express respect for rational nature. Morality thereby requires me to lie in order to express respect for rational nature, thereby overriding any requirement for me to abide by the killer's consent to my project.

the jogger and the murderer merely as means seems absurd. In the case of the killer at the door, the killer's consent is relevant, but any requirement to obtain it is nullified by his violent terms of engagement (and, as a supplement, by our duty to respect humanity), which mandate that we lie. In the case of the jogger, once we have established that consent is required, we must address the issue of competence to consent.⁵⁴ Whenever consent is needed, competence becomes a threshold issue. Clearly, young children, the severely mentally ill, the highly intoxicated and the unconscious jogger are, in their various ways, not competent to consent. In such cases, we have two options, as appropriate. Either we wait until they become competent; or, if waiting is unreasonable, we must turn to a model of hypothetical consent, in which consent is imputed.

v. A blend of actual and possible

If, on the other hand, a person is deemed competent to consent, how do we choose between our various versions of consent? To start, we can disqualify hypothetical consent for the manner in which it hobbles autonomy. Beyond that, one point to note is that several of the issues we have identified are mitigated by the introduction of competence, as we shall see below. A second point is that possible consent, attractive as it is, fails to align with common parlance. Ideally, the conception of consent we are seeking will dovetail with the vernacular. We want a conception of consent that matches, as closely as possible, common understanding of the term. A third point is that neither actual nor possible consent seems to be sufficient on its own. As Formosa writes, we cannot reduce possible consent to actual consent (Formosa, 2017, in press: ch. 3). How then do we choose between actual consent, with all its intuitive appeal, and possible consent, with its theoretical advantages? The answer, I suggest, lies in not making a choice of one over the other. Rather it is to reconcile actual and possible consent, as far as possible. What I propose is that actual consent is required, but that actual consent is defined in terms that *also require* the possibility of consent. Such a conception of actual consent can determine whether someone whose consent is required and who is competent to consent has been used merely as a means. This, then, will be

⁵⁴ How do we determine when consent is or is not the issue without losing ourselves in a confusing mess of circularity? I have addressed this question above, in discussion of permissible noise levels. Further, we might invoke a heuristic: is this person at risk of being treated merely as a means? Generally, it is better that we cast the net too widely, than too narrowly. Even if we catch too many cases, we should nonetheless arrive at the right result.

the conception of consent that I apply to internet privacy in chapter six. Actual consent, to qualify as morally justifying, will necessitate that the person seeking consent also offers the genuine possibility of consent.

Usually, actual consent satisfies the requirement for possible consent. Under normal circumstances, if the local shopkeeper says I can take a packet of jelly beans *gratis*, then I can expect not to be arrested for shoplifting. If an adult explicitly consents to sex, that permission is likely to be morally justifying (unless, say, she was drunk and thus not competent to give consent). As Formosa writes: “Possible consent often (but not always ...) requires actual consent and when we have possible consent then we have a moral authorisation to undertake some action” (Formosa, 2017, in press: ch. 3). Adding in the notion of competence, we can begin to address O’Neill’s three objections to actual consent: that it is unclear precisely what constitutes consent and precisely what has been consented to; that there are instances of hidden coercion; and that problems attend the limited capacity to consent. Each of these issues, I suggest, is remedied to some extent by the application of the notion of competence. Were the Maori chiefs who signed the Treaty of Waitangi competent to consent? Many were illiterate, and we have seen that the very terms of the treaty were unclear, given various oral provisions and promises. What’s more, in cases of limited capacity to consent, the issue of competence is often definitive, as we have already seen.

The issue of competence is also definitive in *some* cases of hidden coercion. If I have been bullied and harassed into sharing compromising pictures of myself, then perhaps that bullying and harassment will have rendered me incompetent to consent. More commonly, however, cases of hidden coercion will not turn on competence, but on exploitation. Does the factory worker freely consent to be a factory worker? What about a sex worker? These people’s consent may well be morally unsatisfactory, but that is not because they are incompetent to consent, like children. Rather, it is because they have no real options. They are constrained to make a choice that they would probably rather avoid. These are difficult cases. No conception of consent, I suggest, will give ready answers. However, the question of competence is a start. Further, notions of relational autonomy and relational privacy I have been advocating recognise that there is significant work to be done addressing underlying inequities, including those that create the

preconditions for exploitative factory work or sex work. This is particularly evident with internet privacy. Online, the issue of competence to consent is crucial for users who are drunk, young, mentally ill or not fully literate. However, issues of exploitation arise too, including in the way services such as Google or Facebook become so popular and dominant that there are hidden pressures to use them, and to share widely with them and through them. Chapter six offers proposals to redress the ensuing inequities of potential exploitation, in the form of both legal and extra-legal measures. In sum, the competence test can bring us some way towards remedying O'Neill's objections, while initiatives that foster relational autonomy, and hence relational privacy, can bring us yet further.

More precisely, what do I mean by actual consent, designed also to incorporate possible consent? As noted above, my scheme is nominally different from, but substantively the same as, Formosa's prescription. Formosa argues that *possible* consent, when required, necessitates actual consent (or the absence of actual dissent) except under three conditions. (1) When they are incompetent to consent. (2) When a rationally required condition means that actual dissent does not revoke an existent authorisation. (3) When a rationally forbidden condition means that actual consent does not grant an authorisation (Formosa, 2017, in press: ch. 3).⁵⁵ Beyond this, I propose three criteria to flesh out a conception. First, the offer of an opportunity to consent must be *bona fide*, so that its focus is the intention of the person seeking the consent, rather than the intention of the person whose consent is being sought. Second, it must take account of the specific personhood of the consenter. And third, genuine possible consent must allow that questions of consent are often iterative and layered, rather than simplistic and singular.

My neighbour is overseas, and I want to chop down a tree that straddles our properties. Fortunately, I know she checks her inbox daily, and so I email her. She spends much of her life online, so I know she will see my query before long. Let's assume also that my neighbour is competent to consent. If my email offers the genuine possibility of consent and my neighbour replies, giving me her actual consent, then I have, it seems, satisfied the requirements of the model I have detailed. What if, however, she doesn't reply? What if my neighbour and I have a frosty relationship? We just don't like one another. In this case, it is quite possible

⁵⁵ The meaning of "rationally required" and "rationally forbidden" is discussed above, in part iv.

that my neighbour would fail to respond purely to frustrate me. In these circumstances, do I need to receive actual consent to satisfy the requirements of possible consent? Is my email sufficient? Usually, silence is not considered to amount to consent. This is a well-established principle in the law of contracts (A.L.C., 1920). However, in 1919, one US state court found that silence maintained for an unreasonable time can amount to acceptance of an offer (A.L.C., 1920: 441). In the context of social work, by contrast, it has been noted that silence can comprise consent, but there is no consensus or clarity as to how and when (Millstein et al., 1994). In short, this is a contested issue. And whereas O'Neill would seemingly conclude that I have given my neighbour the possibility of consent, and hence can chop down the tree, Formosa would conclude that the absence of actual consent means I have no moral authority to chop (Formosa, 2017, in press: ch. 3). As I have signaled, here I depart from O'Neill to align with Formosa. Only actual consent that also affords possible consent is sufficient, unless chopping down the tree is rationally required (perhaps because it is in imminent danger of falling over and killing someone). In other words, there are two requirements: first, actual consent must be obtained; second, the genuine possibility of consent must be offered.

To obtain actual consent that also incorporates possible consent, three criteria must be satisfied. The first criterion is that the offer of consent must be bona fide, and the accompanying moral test hinges on the intention of the person seeking the consent, rather than the consenter. If I am seeking consent, what has been my action? What has been my intention? One heuristic is to ask: have I offered the possibility of consent *to the best of my knowledge*? What we are seeking, paradoxically, is *actual* possible consent. This prohibits willful ignorance on the part of the person seeking consent. In this case, then, even if my neighbour gives me actual consent, this consent will not be morally justifying unless, for instance, my proposed action has been described in such a way as to spell out its salient features. In this way, my conception of actual consent steps away from a *caveat emptor* approach, in which the consenter must be diligent in giving or withholding consent. The model I propose, by contrast, invokes *caveat venditor*. In this model, it is the seller (that is, the person proposing the action) and not the buyer (the person who will be affected by the action) who has the moral responsibility to ensure that consent is not merely a formality, but morally justifying. This is

particularly relevant on the internet, where data use and re-use is complex and unpredictable. The onus is on companies, for instance, to obtain the actual consent of users, and actual consent must also contain the genuine possibility of consent.

The second criterion is that possible consent must take into account the particularities of the person whose consent is being sought. If my neighbour is a fanatic greenthumb who curates every leaf and bloom with forensic precision, then the onus upon me for my tree-felling proposal is considerably higher than if she were botanophobic. If a woman admits to a low tolerance for alcohol and then, after two drinks, begins slurring and stumbling, it will be difficult to offer the possibility of consent for sex. As O'Neill writes, taking account of the specific personhood of the other is a key component of acting in accordance with the FH:

An adequate understanding of what it is to treat others as persons must view them not abstractly as possibly consenting adults, but as particular men and women with limited and determinate capacities to understand or to consent to proposals for action (O'Neill, 1989: 105).

It matters if my proposed action will affect a stranger, a barely-known colleague, my close friend or my wife. It matters because in each case our relationship is different, and in each case what I know about this other is relevant. I might know, for instance, that one particular colleague is always chirpy on Friday afternoons. If I have a potential project that affects her significantly, would it do to wait until then to seek her consent? Perhaps. However, if I know that her Friday high is caused by an untreated personality disorder, or a weekly martini-infused lunch, then arguably her capacity is limited on Friday afternoons in a manner that renders her incompetent to consent. We must take account of the fact that we know someone is hard of hearing, or struggles to comprehend English, or is in the manic phase of a manic depression. These factors have the potential to compromise, and perhaps invalidate, their consent. To avoid treating another as a mere means, I can only offer the possibility of consent while bearing in mind her quirks and eccentricities, strengths and weaknesses. Again, this is relevant online, where consent must be offered in a manner appropriate to users, be they children, or internet novices, or those with poor literacy.

Finally, a third criterion is that genuine consent must be iterative and layered as required. Consent, to be morally justifying, tends to be an ongoing process, rather

than a dialogue box that needs to be ticked once only. For a sexual encounter, consent is rarely given in the form of an explicit answer to a straightforward proposal. Rather, consent takes the form of subtle cues, implicit signals and explicit declarations, all of which can then be revoked with any one of a number of declarations of dissent. After obtaining an initial consent, I shouldn't then insert earplugs in order to avoid hearing a subsequent reversal. Here, I may have actual consent, but need to *keep offering* the possibility of consent. In many cases, as an individual begins to understand more about any given situation, the issue of consent needs to be revisited. Once again, this point is evident in our online interactions. Google emerged as an advertising-free search engine; it now offers email (Gmail), maps (Google Maps) and cloud-based document services (Google Docs, Google Drive), while engaging in sophisticated personalised advertising that relies on secret algorithms and users' browsing histories (see chapter two). For Google to offer the genuine possibility of consent it must enable its users to keep re-consenting as its services change, and as its use of individuals' data changes. Online, re-consent is required. What's more, issues of consent are often layered. If I am reading the *New York Times* inside Facebook's Instant Articles application, both companies' privacy provisions and consent policies will be relevant, even as they interlock in ways that are difficult to discern and understand. The issue of consent can thus involve several parties (see chapter six).

My model of individual competence and consent can be summarised as follows:

1. Is consent required?
 - a. If no, look elsewhere for ethical guidance.
 - b. If yes, proceed to question 2.
2. Do any rationally required or forbidden conditions override consent?
 - a. If yes, they render consent and dissent irrelevant.
 - b. If no, proceed to question 3.
3. Is the person competent to consent?
 - a. If no, hypothetical consent is required (unless we can wait until the person becomes competent, in which case we can then proceed straight to b).
 - b. If yes, actual consent is required.

- c. Actual consent *also* requires the offer of the possibility of consent, which includes:
 - i. A focus on the intention of the person seeking consent;
 - ii. Taking into account the vulnerabilities and specificities of the person whose consent is sought; and
 - iii. A recognition that consent ought to be adequately iterative and layered.

With my Doisneau homage, *Kiss at the Sydney Opera House*, I have inadvertently created a masterpiece, but what consent do I need? First, I must ask whether consent is required. Will the couple in the image be significantly affected by my action? Given that I am planning to exhibit and sell the image, the answer is yes. After satisfying ourselves on the second question that there are no rationally required or forbidden conditions, the third question then becomes whether the couple is competent to consent. Are they drunk? Are they children? Are they mentally impaired? If they are incompetent, there are two possible courses of action: either I wait until they become competent; or someone competent consents on their behalf. Circumstances will determine which of these options is the right one. If they are drunk, I can seek their consent later. If they are mentally impaired, or they are children, then hypothetical consent is required, in the form of a guardian or parent who can decide on their behalf. Sometimes, incompetence may warrant deception or coercion, as in the case of the unconscious jogger, but clearly this is not the case here.

By contrast, if they are competent, I must obtain their actual consent. In doing so, I must also offer them the genuine possibility of consent. This involves satisfying the three criteria I have just elaborated: that consent be a bona fide case of *caveat venditor*, with a focus on my intention; that consent take account of the specificities of the couple; and that consent be iterative and layered, as appropriate. If I am a famous photographer with a reputation for bullying, or if I omit to mention that I plan to exhibit and sell, then I am probably failing to satisfy the first criterion. If the couple clearly have only a limited grasp of English, or are from a culture where it is impolite to say no, then I may well be failing criterion two. And if initially I had planned only to exhibit in a gallery, but now I plan to use the image on T-shirts and TV commercials, then I am required to have the

couple re-consent. None of this addresses in detail several hard issues, such as: whether their silence can ever be construed as consent; and whether I can proceed even if I cannot identify and locate them. To resolve such issues, the law can, and should, provide binding guidelines. The law sits atop individual consent. Indeed, the law, as an expression of collective consent, can and should provide guidance for many of the challenging issues I have raised, as I explore below. In what cases is consent required? When is a person considered competent to consent? What is to be the agreed definition of consent in cases of competence? And how do we define hypothetical consent in cases of incompetence?

My aim has been to spell out a clear and effective conception of individual consent in order to apply Kant's formula of humanity to internet privacy. As we have noted, consent does not satisfy every case where we seek to apply the mere means principle. Nonetheless, it is *one* powerful tool by which we can apply Kant's formula. Once we have established that a person is competent to consent to an action that significantly affects them, consent is required. Specifically, what is required is *actual* consent, and this actual consent must also be defined with reference to the *possibility* of consent, as described above. In this way actual consent can best counter the opacity of intentionality, accommodate the complexity of human engagements and also allow for the frailty, vulnerability and irrationality of individuals. Such an approach would also align with common parlance. Legally and ethically, what we should be required to seek is, paradoxically, *actual possible consent*. This will be consent that adopts various merits of both actual and possible consent, while dispensing with some of their weaknesses.

Meanwhile, we must simultaneously broaden our focus to recognise and redress some of the factors that work to stifle consent and competence. For instance, we ought to identify and ameliorate the hidden and systemic coercion that operates in certain contexts as well as work to counter the institutional forces that tend to clip people's ability to be self-determining, self-governing and self-authorising. Herein lies an acknowledgement that there are certain preconditions for autonomy, as there are for competence and consent, and these must be fostered. For Mackenzie, for instance, an agent cannot be self-determining unless certain structural sociorelational conditions are in place. If agents stand in relations of

subordination, subservience, deference, or economic or psychological dependence, they are unlikely to be autonomous, even if the agents themselves endorse these subordinate, subservient, or dependent positions (Mackenzie, 2014). Autonomy is relational, and so too privacy. Our competence to consent, and our consent itself, is intertwined with our socially-constituted identities, and with our relationships. Only once society becomes more fair and equitable can consent and competence become more comprehensively morally justifying. In the meantime, actual possible consent can help, particularly when combined with the collective consent of the law.

IV - Collective consent

Stephen Gough is a former marine whose abiding love of public nudity and long walks have earned him the nickname the Naked Rambler. In 2003-2004, Gough walked the length of Great Britain wearing only boots, socks and a hat. Two years later, while repeating his cross-country trek, he was arrested. Ever since, he has spent most of his time in prison. His life now follows a rhythm as regular as the seasons: public nudity, arrest, court, jail; public nudity, arrest, court, jail; and so on. Gough argues that it is his human right to be naked in public. Indeed, under British law, public nudity is not a crime. However, authorities have imposed an Anti-Social Behaviour Order, or “Asbo”, making it unlawful for Gough to be naked in public (Miller, 2015). As one barrister noted: “The result is that the only person in the country who actually wants to wander naked around the streets of Winchester is also the only man in the country who commits a crime by doing so.” Had he crossed the English Channel, the story may have been different. On the Continent, Gough’s nudity may well have been legal (Whitman, 2004: 1196-1197). My point here has nothing to do with the rightness or wrongness of public nudity, and hence whether or not it should be legal. Rather, my point is that Gough himself does not have sole say in determining where the limits of his own privacy are to be drawn. Willingly, Gough has brought his naked body, the acme of what is usually considered private, out into the public sphere. He has consented to his own public exposure, and has thereby implicitly consented to others looking at his naked body. Through his actions, he has sought to shrink his condition of privacy and to give up some of his right to privacy. Unfortunately for him,

however, UK legislators, regulators and judges have decided that he is not entitled to do so. On the limits of his personal privacy, Gough's individual determination has been overridden by a collective pronouncement.

In the previous section, I sought to build a model for applying the formula of humanity by articulating a robust conception of individual consent. However, sometimes individual consent isn't enough. Even if the conception of consent I have given is the very best conception of individual consent, there is more work to be done. Throughout this thesis, I have been talking about individual privacy. It might seem logical that the limits of privacy ought to be set by the exercise of individual choice, as expressed by consent and dissent. However, as the case of the Naked Rambler reveals, the limits of privacy are not always solely the matter of the individual whose privacy is in question. This is the same point I made in chapter four: that the decision of which swimsuit I can wear is not entirely up to me. There are other prescriptions that sit over and above individual standards. These prescriptions include laws, comprising Kant's "united will of the people", which can work to qualify, buttress, amend, invalidate or otherwise affect individual consent. These laws amount to "collective consent". To put it another way, there are two important sources of authority that are independent of our actual will or consent: our inherent dignity; and the united will of the people (Formosa, 2017, in press: ch. 3). The former prevails in the case of the killer at the door and the case of the unconscious jogger; the latter is the subject of this section.

On the internet, collective consent has an especially vital role to play. In March 2015, following the bankruptcy of US retail chain RadioShack, the company's assets went to auction. These assets included the personal data of its customers. In all, the company put up for sale 13 million e-mail addresses and 65 million customer names and physical address files (Brustein, 2015). Presumably RadioShack's notice-and-consent provisions were inadequate to cover such a sale. In chapter three, I described Helen Nissenbaum's transparency paradox: either notice-and-consent provisions are overly simple and inadequate; or they are overly complicated and hence mostly unread (Nissenbaum, 2011: 36). Studies have shown that many internet users have blindly agreed to terms that include giving up their first born and selling their soul (Obar and Oeldorf-Hirsch, 2016;

Smith, 2010). O'Neill talks of the opacity of intentionality that plagues consent; nowhere is this more evident than online. Studies consistently show that privacy policies are misnamed, unread, unreadable and incomprehensible (Cranor and Reidenberg, 2002; Turow et al., 2007: 723-724; Nissenbaum, 2011: 35-36). Moreover, research has shown that self-regulatory bodies set up in the US to enforce website compliance with notice-and-consent provisions are sometimes ineffectual (Komanduri et al., 2011).

As we turn from individual consent to collective consent, our focus becomes the *right* to privacy. As such, our focus shifts from virtue to justice, from ethics to social and political philosophy. This is challenging. As Kant wrote, defining “right” is sufficiently difficult that it “might well embarrass the jurist” (Kant, 1996a: 6:230). However, the shift is apt. When tackling the issue of internet privacy, we *must* contemplate both virtue and justice (or right).⁵⁶ As O'Neill writes: “I suspect that ... failure to think about justice and virtue in tandem is likely to lead to blinkered and ungenerous, as well as implausible, visions of life, action and politics” (O'Neill, 1996: 6). Privacy, I have been arguing, is an individual concern, but also a social concern. Privacy is a matter that impacts individual well-being, but also the social fabric. It can only ever impact the individual as a being-in-relation. It is a public good and a private good. What I hope to render is an account of social and political philosophy that is consistent with the ethics I have described, and which then enables us to build a more comprehensive prescriptive framework for privacy. This also takes into account Kant's view of rights as highly significant. For realist political philosophers such as Machiavelli and Carl Schmitt, the political realm should not be bound by concepts that can be traced back to good and evil (see Formosa, 2008: 158). Kant, by contrast, argued that politicians are invariably answerable to standards of right or justice in their exercise of public duties (Wood, 2014: 76). For Kant, “all politics must bend its knee before right”; and “right must never be accommodated to politics, but politics must always be accommodated to right” (Kant, 1996c: 8:380; Kant, 1996b: 8:249; quoted in Formosa, 2008: 157). The state, after all, has powers of coercion.

⁵⁶ In a linguistic ambiguity that continues to trouble translators, the German noun “Recht” means both “right” and “justice/law”, and the German adjective “recht” means both “just” and “right”.

Kant argued that persons, as self-legislating members of the realm of ends, are subject to two kinds of constraints: internal and external. Internal constraints, comprising the moral law that all persons legislate for themselves, are required by the categorical imperative. This internal legislation is then supplemented by external legislation enacted by the state (Kant, 1996a: 6:220). Answerable to standards of right, external legislation must not contradict the moral law. Hence the law cannot legitimately implement slavery, or apartheid. For a law to be legitimate, it must be just (see discussion below). For Kant, both virtues and rights are thus concerned with morality, but while *virtues* pertain to the morality of individuals and their behaviour, *rights* are what individuals are given by the state to protect them, *inter alia*, against others acting in a way that improperly impinges upon their autonomy. The only natural right is that of freedom itself:

There is only one innate right. Freedom (independence from being constrained by another's choice), insofar as it can coexist with the freedom of every other in accordance with a universal law, is the only original right belonging to every man by virtue of his humanity (Kant, 1996a: 6:237).

Right, for Kant, is the sum of the conditions under which the choice of one can be united with the choice of another, in accordance with the universal law of freedom (Kant, 1996a: 6:230).

Freedom, however, cannot be unfettered. If the freedom of one is to coexist with the freedom of every other, constraint is required. This is the constraint provided by internal and external legislation. At the heart of Kant's politics, then, sits a paradox: freedom requires constraint. According to the formula of humanity, as well as other formulations of the categorical imperative, there are limits to the ways in which I ought to behave, just as identical limits prescribe the ways in which you ought to behave. The formula commands that I, you and every other rational creature ought never to treat another as a mere means, but always as an end in herself. Hence I am duty-bound to uphold the formula, which I can regard myself as self-legislating. Meanwhile, the state has, and should have, the power to coerce individuals to behave in a manner that aligns with the formula in the interests of promoting individual freedom. State coercion should thus have as its highest goal the promulgation of an order wherein justice – and hence mutual freedom – prevails. As Kant wrote in 1797,

Coercion is a hindrance or resistance to freedom. Therefore, if a certain use of freedom is itself a hindrance to freedom in accordance with universal laws (ie, wrong), coercion that is opposed to this (as a *hindering of a hindrance to freedom*) is consistent with freedom in accordance with universal laws, that is, it is right (Kant, 1996a: 6:231).

In this way, laws that embody the united will of the people can rightfully coerce citizens to behave in some ways and not in other ways.

In Kant's writings, the notion of "general will" or "united will" recurs often. In "Theory and Practice", Kant writes that the sovereign must recognise that he is obliged by the social contract to "give his laws in such a way that they could have arisen from the united will of a whole people and to regard each subject, insofar as he wants to be a citizen, as if he has joined in voting for such a will" (8:297, quoted in Rauscher, 2016). The general will can, for instance, extend to revenue raising. As Kant writes, the commander-in-chief (*Oberbefehlshaber*) can levy taxes to provide for the poor and to fund orphanages. People submit to the state willingly, Kant writes, to help look after those unable to look after themselves (Rauscher, 2016). Property rights are also an expression of the common will, given that they comprise "appropriation (*appropriatio*) as the act of a general will (in idea) given an external law through which everyone is bound to agree with my choice" (Kant, 1996a: 6:259). Such general will can be regarded, to use Wood's phrase, as omnilateral consent: "Even private right (the right of individual property) depends, as a peremptory right, on *omnilateral* consent ..." (Wood, 2014: 77). It is in the law, I am arguing, that we can locate a second layer of consent that overarches the first layer of individual consent. The law, when legitimate, operates as an expression of *collective* consent. And this collective consent can work to qualify, buttress, amend, invalidate or otherwise affect individual consent.

However, not all laws are legitimate. Only a law that is *just* can qualify as the united will of the people. I have already noted that morals trumps politics, quoting Kant's dictum that "all politics must bend its knee before right". Following Montesquieu, Kant wrote that the "general united will" consists of three persons: legislator; executive; and judiciary (Kant, 1996a: 6:313). Of these, it is the legislator who can embody the "concurring and united will of all" (Kant, 1996a: 6:314). What does he mean thereby? Is he proposing that the legislator requires the *actual* consent of every citizen for every law? That would be impossible.

Rather, as with individual consent, Kant returns to the notion of *possible* consent, writing in “Theory and Practice” that no law may be promulgated that “a whole people could not possibly give its consent to” (8:297, quoted in Rauscher, 2016). Instead of empirically gleaned consent, Kant advocates a *rational possible unanimity*. For example, a law would be unjust, Kant argues, if it provided hereditary privileges only to a certain class of subjects. As dignified rational beings, all persons have an equal moral standing; but such a law suggests that some have less worth than others due to birth, and then seeks to perpetuate such inequality. Hence those excluded from such privileges could not possibly assent to it (Rauscher, 2016). However, a law imposing a war tax could be just, even if many citizens openly disagree with it, as such a tax does not assail the irreducible dignity of all persons (Rauscher, 2016). A law, merely by its enactment, does not necessarily comprise the united will of the people. A law is not just simply because it is a law. *Prima facie*, it may be fair to assume that a law is a just expression of the collective consent. This assumption may, however, be challenged and overturned. At times, the legislature will make mistakes, passing laws to which the populace could not possibly have assented. These laws are unjust and illegitimate.

For internet privacy, the issue of what constitutes a just law arises most obviously in the case of blanket surveillance by government agencies (see chapter two). Can legislation that authorises the NSA to conduct blanket surveillance of all US citizens be considered a just law? To answer, Kant would ask whether the law was one the whole populace could possibly have assented to. The fact that it might adversely affect people or be unpopular, like a war tax, is of little consequence. However, if a law attacks the dignity and equal moral standing of citizens, then it is unjust. Hence if a surveillance law effectively privileges the rights of one group of people over another – those, say, of a particular race, or religion, or socio-economic status - then it is necessarily unjust. What’s more, if a law is *indiscriminately* a violation of the dignity of humanity, it would be unjust. As well, there are also Kantian arguments that governments have a duty to openness. Where this duty prevails, and yet where laws are drafted and executed in secret, there is a greater chance that a law is unjust. For Kant, legislators must abide by principles of publicity (Wood, 2014). Further, what about laws authorising the NSA’s mass surveillance of *non-domestic* citizens? The international nature of the

internet enables global surveillance, which significantly complicates the question of possible consent. I return to the issue of the justness of laws authorising blanket surveillance in the next chapter.

I opened this section with the case of the Naked Rambler. Let us presume that Gough's naked rambling is neither forbidden nor commanded by the moral law. However, external legislation has forbidden his behaviour. As such, Gough has spent most of the past decade in jail. During his naked ramblings, Gough flew a white flag from his backpack bearing the hand-scrawled slogan, "The Freedom to be Yourself." Kant's formula protects the freedom to be yourself. But it also limits that freedom, by taking account of the freedoms of others. Freedoms are offset by obligations (to respect others' freedoms). Rights are offset by duties (to respect others' rights). Gough's freedom to be himself is offset by the collective right that privates stay private. Of course, there is an argument to say that the laws and regulations by which Gough has been imprisoned are unjust. If, however, we assume those laws and regulations are just, then, in the matter of setting aside Gough's bodily privacy, *collective dissent overrides individual consent*.

Consider a second case, this time hypothetical. While on holiday in Melbourne, I am dining with my wife in a suburban restaurant, which is almost empty. The only other party comprises two men deep in conversation, *sotto voce*. Due to a quirk of acoustics, however, my wife and I can overhear every word. Normally, I would feel obliged to respect this party's privacy. In a restaurant, my usual assumption would be that I should not eavesdrop. In our scenario, I would suppose that I should respect the men's privacy. If, however, I overhear these two men talk of hosting a child pornography website, then I am committing a criminal offence under Victorian state law if I fail to report that conversation to police (Victoria, 2016). *Prima facie*, the men in conversation are entitled to privacy. If I want to pass on the contents of their conversation, it seems, I would need to obtain their consent, or else I would be using them merely as a means. However, under amendments to the *Victorian Crimes Act* made in 2014, the law requires me to report their conversation, or else I face a maximum penalty of three years in prison. Once again, the individual limits of privacy have been overridden by a collective pronouncement. These two individuals never consented for their privacy to be encroached upon, but the law rendered such an encroachment

obligatory. In this case, when it comes to sharing details of an overheard conversation, *collective consent overrides individual dissent (or, more specifically, absence of required consent)*.

Collective consent overarches individual consent. I may wish to keep my financial dealings entirely to myself, but am obliged by law to share those details with the tax office. I may prefer to be naked when I swim at the beach, but the law obliges me to cover up. This is evident on the internet too. Laws in several jurisdictions aim to bolster individual consent in the face of spyware and revenge porn, and to punish infractions of such consent; laws prescribe the individual consent that companies are required to obtain before they can share and sell user information; and laws authorise the Australian government to store and monitor my emails in the fight against terrorism, whether or not I consent to that monitoring (Australian Government, 2017: the above laws are discussed in chapters three and six). As long as these laws are just, they are legitimate. And for now I merely wish to make the larger point that whenever there is collective consent, that collective consent *trumps* individual consent. When it forbids slavery or mandates a random breathalyser test of a driver's blood alcohol reading, collective consent *nullifies* individual consent. When it prescribes that cars may only be driven by those over the age of 17, collective consent *specifies* who may consent to drive, and who may not. When it provides that individuals may not trespass without the owner's consent, collective consent *reinforces* individual consent. Kant's "united will of the people" thus stands astride individual consent, overarching and overseeing it by limiting, buttressing, qualifying, nullifying, specifying and otherwise affecting individual consent. By contrast, whenever collective consent is silent, individual consent reigns. As long as these laws are just, then such collective consent is entirely warranted. Simply, individual consent must be tempered by justice, and justice is contained in collective consent. It is in collective consent that we find the *right* to privacy.

In privacy law, collective consent can involve protecting the public interest, underscoring that privacy is as much a social concern as an individual concern. In 2014, the Australian Law Reform Commission recommended the introduction of a cause of action for serious invasions of privacy, but only those invasions "that cannot be justified in the public interest ..." (ALRC, 2014). As the ALRC wrote:

“A plaintiff should not be able to claim that a wrong has been committed - that their privacy has been seriously invaded - where there are strong public interest grounds justifying the invasion of privacy.” Encapsulated in the notion of the public interest is the idea of a balancing of rights, interests and values. Indeed, just as the law prescribes and protects the right to privacy, so too it prescribes and protects other rights. The bulk collection of health information may lead to better treatments, but may compromise individual privacy. The mass surveillance of citizens may increase citizens’ security, but *does* compromise citizens’ privacy. As I shall explore in more detail in the next chapter, the right to privacy must be balanced against other rights, obligations and freedoms, and it is the collective consent of the law that must spell out how that balance is to be struck.

For Kant, it is in the law that our rights, including the right to privacy, are expressed and protected. Ethical precepts and social norms can give us the expectation of privacy, and can even help to provide and protect the condition of privacy, but only the law can prescribe the right to privacy. We can think of these state-legislated protections of the right to privacy as expressions of collective consent. In other words, individual consent will enable us to apply the mere means principle of the formula of humanity; but the formula can only be applied within the framework of the law, which must itself be in harmony with the formula and the moral law. That law (as long as it is just) can be thought of as a collective consent that overarches individual consent, and which gives us our right to privacy.

V – How two layers of consent mesh with control and access

In chapter three, I argued for an access model of privacy. Following Gavison (1980), Reiman (2004) and others, I argued that privacy is always about a restriction on access to ourselves. Sometimes these restrictions can involve our own control. In such cases, our privacy is determined by our control and its value lies in this control. Here, our right to privacy denotes our right to set the limits of our privacy. At other times, however, these restrictions will be the result of externally-imposed limitations on access to us and information about us. In such cases, our privacy is set by external forces, such as the law and social norms. And

it is in the law that we find our right to privacy, which gives us the right to ensure that these externally-imposed limitations are observed. In this final section, I want to make a fairly straightforward proposition: that individual consent equates to control; and that collective consent equates to externally-imposed restrictions upon access. Hence, just as neither control nor externally-imposed limits on access is sufficient to define privacy, so too neither individual consent nor collective consent is sufficient to determine an individual's privacy. Control must be complemented by externally-imposed restrictions on access, just as individual consent must be complemented by collective consent. Where there is a clash, access/collective consent ought to triumph (so long as it is just, and does in fact express the united will of the people).

Consider the act of getting dressed in your bedroom, an act usually regarded as private. If Tom secretly watches Cassidy via webcam, he is presumably using her merely as a means and thus violating her privacy. If, on the other hand, Cassidy signs a form granting blanket permission, it would seem that Tom has not treated her merely as a means. The issue of individual consent is decisive. However, there is a further consideration which transcends the issue of individual consent. That is: does a law apply? We can imagine, for example, a law forbidding such surreptitious use of cameras, regardless of individual consent. In such a case, the law would dictate that the condition of privacy must be respected. A legal right to privacy would prevail, even if an individual would prefer to waive it. Conversely, one can imagine that in certain cases (as when there exists a well-founded suspicion that someone is plotting a serious crime), that surreptitious filming would be permitted under the law, irrespective of dissent. In this way, collective consent overarches individual consent. Together, individual and collective consent can be definitive in determining whether or not there has been a breach of the mere means principle and hence an unacceptable intrusion onto privacy, taking account not just of an individual's wishes, but of state-imposed limits on permissible behaviour. This two-tier model of consent can help to determine when the condition of privacy ought to prevail, and when a right to privacy ought to exist.

The privacy settings of Twitter are an example. To set up a Twitter account, users first agree to stipulated terms and conditions, which they can then tailor to be

more open or less open. They can, for example, specify whether or not other users can tag them in photographs (Twitter, 2017).⁵⁷ Therein lies control. This control is sometimes more illusory than actual, as Twitter admits: “Tip: What you say on the Twitter Services may be viewed all around the world instantly. You are what you Tweet!” Nonetheless, users have *some* control, enacted by adjusting their settings. Meanwhile, this control is overseen by a long list of laws. In Australia, the *Privacy Act 1988* puts limits on the way a company such as Twitter can share “personal information”, and even stricter limits on “sensitive personal information”. Moreover, laws exist to prevent the sharing of explicit images. This is most obvious with content depicting children. In Australia, it is a crime for anyone under the age of 18 to “sex”, and it is a crime for anyone else to share such images. Individual consent and control are rendered irrelevant; instead, the law limits access by prescribing that such private images are not to be shared, irrespective of consent. Hence 16-year-olds potentially face jail for exchanging nude selfies (Lawstuff, 2016). Other laws limit adults’ ability to consent. And, as noted above, a 2015 law enables the Australian government to access the electronic metadata of citizens for two years, irrespective of citizen consent (AustralianGovernment, 2017).⁵⁸ Privacy on Twitter thus involves both the control expressed in individual consent and the significant externally-imposed limitations on access set by the collective consent of the law.

Now let’s return to two thought experiments first raised in chapter three: Judith Jarvis Thomson’s X-ray device and Adam Moore’s DNA residue. Further, let us relocate these cases to the internet. Hence we can transform the X-ray device into a type of spyware which secretly monitors our every keystroke and online activity, thereby recording every website we visit, every email we send and every picture we post. When this spyware installs itself on our browser, we are totally oblivious. As noted previously, such technology exists (Nissenbaum, 2010: 21-67; Schneier, 2015: 62-77). Let’s call this the Spyware scenario. Clearly, on this scenario, we have lost control. Our ability to grant or withhold consent has been rendered useless. Given that such malware has the ability to track our every digital interaction surreptitiously, our condition of privacy has been compromised.

⁵⁷ Tellingly, the default setting allows users to be tagged by others. This is typical of social media, where default settings tend to encourage sharing, not privacy. In real life, by contrast, the default position tends to be for privacy, not sharing. See chapter six.

⁵⁸ See chapters three and six for more detail on Australia’s *Privacy Act* and on sexting laws.

But has there been a breach of our right to privacy? That depends on the law. If such spyware is being used by a government agency, and if the agency is authorised to use such technology by a just law, then there is no infraction of my right to privacy. If, however, it is an individual or company installing such malware, and laws exist to prevent unauthorised access to individuals' computers, then my right to privacy has been violated. As I discuss in chapter six, such laws exist. Again, the collective consent of the law overarches individual consent. In this case, as in the case of revenge porn, laws exist to criminalise the failure to respect individual consent. Here collective consent works to reinstate the power of individual consent and user control.

In the age of the internet, Thomson's thought experiment becomes more plausible and pressing. So too Adam Moore's DNA scenario. If I walk in the park, someone could potentially collect an abandoned strand of hair and use it to map my DNA, without my consent, or even knowledge, before uploading that information to the web. Indeed, DNA databases are proliferating, both among police and privately-owned companies geared towards people interested in genealogy, generating heated debate about appropriate ethical responses (Chadwick and Berg, 2001). Clearly, I have no control over what happens to my errant follicle. Consent is meaningless. However, it seems that my privacy would be invaded if someone collected my hair and mapped my DNA. The act of uploading that information to a publicly available database would be a further and more flagrant violation. Given that control and consent have been invalidated, collective consent becomes key. For instance, a law could make it illegal to collect DNA without consent. Such a law was proposed by the Australian Law Reform Commission in 2003, which argued DNA collection should be criminalised if it occurs "without the consent of the person concerned or other lawful authority" (ALRC, 2003). Alternatively, laws could make the collection of DNA without consent legal in some circumstances. In the US, for instance, US courts have upheld laws giving police broad powers regarding crime suspects, including the finding that police could lawfully collect a suspect's DNA without his consent and without a warrant (Lynch, 2015). A third option is to make the collection of DNA mandatory, irrespective of consent. In Kuwait, following an Islamic State suicide bombing in July 2015, parliament legislated to make DNA testing mandatory for all 4 million citizens and foreign residents (AlJazeera, 2015). Given the inefficacy of

individual consent, and hence control, collective consent in the form of just law has a significant role to play in specifying limitations upon access to my genetic code, with all its profound insights into who I am, and also who my family is.

Individual consent enables control. With consent, users can control access to themselves and to information about themselves. Online, however, consent and control are often problematic. Hence *collective consent in the shape of the law* has a significant role to play. By imposing restrictions upon access, the law can establish and protect the right to privacy. It can do so by reinstating the validity and value of individual consent; alternatively, it can do so by setting the limits of privacy, irrespective of individual consent. In this way, individual consent can be regarded as a rule of virtue. Collective consent, by contrast, is the rule of right. As Formosa notes, it is only when the coerced rule of right is coupled with the noncoerced rule of virtue that “humanity’s end of an enlightened age of peace” can be realised (Formosa, 2008: 181).

Conclusion

At the end of *The Dark Knight*, Lucius Fox types in his name, and the surveillance apparatus designed by Wayne Enterprises self-destructs. “Sometimes people deserve to have their faith rewarded,” Batman mumbles. The Joker has been defeated; innocent lives have been saved; the surveillance system has been detonated. Of course, that’s not to say Batman, or someone with fewer scruples, will build another just like it. This is no uncomplicated happy ending. The ethics and politics of these scenes are highly ambiguous (CriticalCommons, 2016). Lucius and Batman have jettisoned some ethical principles and civil rights to prevent a calamity. First, Lucius condemns the surveillance apparatus and its power; then he and Batman use it to thwart the Joker’s terrorist plot. Are they right to do so?

In this chapter, I have argued that an application of Kant’s formula of humanity, which exhorts us never to use someone merely as a means, can illuminate the ethics of internet privacy. To do so, we must obey (in general, at least) the consent and dissent of those significantly affected by our actions. However, consent can

be fraught, especially online. As such, we need to start by asking: is the person affected competent to consent? If yes, the test to apply is one of actual consent, defined also to incorporate elements of possible consent. If no, then consent must be scaffolded, and hypothetical consent comes into play. That is, consent must be given or withheld on someone's behalf. Even so, individual consent does not settle the matter. For a thorough application of the formula of humanity, we must then shift from ethics to politics, from virtue to right, and thus think also of collective consent, in the form of Kant's "united will of the people", as expressed in just laws. These laws give us the right to privacy, and help to balance privacy against other rights. This collective consent can reinstate the power of individual consent, it can render individual consent null and void, or it can specify when individual consent does or does not apply. In short, collective consent overarches individual consent, working to limit, qualify, buttress and otherwise affect it. In this way, I have aligned my access conception of privacy with a two-tier account of consent: control equates to individual consent; externally-imposed restrictions upon access equate to the collective consent of just laws.

With their elaborate surveillance apparatus, Lucius and Batman were ostensibly using the citizens of Gotham merely as a means to capture the Joker. On the model I have described, were they entitled to spy? *Prima facie*, no. Seemingly, they were not entitled to engage in such surveillance as they were not authorised by a just law. Instead, they were acting as vigilantes, imposing an *ad hoc*, extra-judicial justice. They failed to obtain the individual consent of citizens, and were not authorised by collective consent. However, a closer analysis reveals that in fact Batman and Lucius *were* authorised to spy. On the model I have described, the first question asks whether consent is required. Clearly, the consent of residents is needed for such surveillance. The second question then asks whether there is a rationally required or forbidden condition. In this case, there is: Batman is *rationally required* to save the citizens on the ferries, whom the Joker is about to kill. The moral law *demand*s that Batman spy. The requirement for consent is nullified by the requirement to save lives that are in imminent danger. Rather than using the citizens merely as means, he is expressing respect for humanity by spying on all of them to save some of them. A government would be similarly entitled in the face of an *imminent* threat. In the face of a more *generalised* threat such as terrorism, however, a government agency such as the NSA can only be

justified in analogous surveillance if it is authorised to do so by a law to which all citizens could possibly have assented. What's more, this holds only for domestic surveillance. When such surveillance is conducted internationally, problems arise. Just how can we obtain global possible consent? I return to this issue, and to the requirement that governments act in accordance with Kant's principles of "publicity", in the next chapter.

In the absence of Batman, we can use *actual possible* consent as a test in order to avoid using people merely as means. If someone finds a strand of my hair in the Vondelpark, they ought to seek my consent before mapping my DNA and uploading that information to an internet database for all the world to access. However, people don't always do as they ought. Hence collective consent, in the form of external limitations upon access imposed by the law, has a role to play. Indeed, the law's approach, here and elsewhere, might well be to reinstate my autonomy by requiring my consent. In this way the law would unequivocally be bending its knee before right, and before virtue.

Chapter 6

A privacy-respecting *cosmopolis*

The Social Network is a film about Mark Zuckerberg, the wunderkind coder behind Facebook. It opens in 2003, as the 19-year-old Harvard undergrad is being dumped by his girlfriend, Erica Albright. Humiliated, Zuckerberg returns to his dorm, where he writes an offensive blog post about her before, in a frenzy of drunken coding, hacking into college databases to steal photos of female students. Using algorithms, he then builds a website called Facemash, inviting users to rate the attractiveness of students. Now it's Albright's turn to be enraged. "The internet's not written in pencil, Mark, it's written in ink," she says. "And you published that Erica Albright was a bitch, right before you made some ignorant crack about my family's name, my bra size, and then rated women based on their hotness." Zuckerberg is stung, but Facemash is a hit, becoming so popular so fast that it crashes a part of Harvard's computer network. Admittedly, *The Social Network* is a feature film; but the key plot details are true to life. Zuckerberg had been jilted by a woman; he did use his hacking skills to copy digital profile pictures; and he did build Facemash for users to compare the "hotness" of female students (having first abandoned the idea of comparing them to farm animals) (Hoffman, 2010). In the film as in real life, Zuckerberg was reprimanded and Facemash was closed. Nonetheless, Facemash was the prototype for Facebook, which Zuckerberg launched the following year. Like Facemash, Facebook allows users to pore over the lives of others and soon had privacy alarm bells ringing. Even so, Facebook proved so successful that even Zuckerberg was surprised. As he later said, "People are more voyeuristic than what I would have thought" (Hoffman, 2010).

Once we accept that privacy matters, and once we further accept that our internet use is challenging privacy, the issue becomes: what can we do about it? This chapter is an attempt to offer practical solutions, in the form of legal and extra-legal measures that both implement and complement the theory outlined thus far. In previous chapters, I have described Panopticon 2.0, in which a triple threat to privacy exists from individuals, corporations and governments. I further argued that privacy is about restrictions on access, and that these restrictions sometimes

take the form of control, before showing that privacy matters for reasons including dignity, autonomy and relationships. Then, in chapter five, I drew on Kant's formula of humanity to argue that privacy can be illuminated by a two-tier model of consent, in which individual consent is supplemented by the collective consent of the law. In this chapter, I invoke these descriptions and prescriptions in order to apply them. In the first section, I give a general overview of practical solutions via a return to the triple challenge established in chapter two. One promising solution, I suggest, lies in tort law. In the second section, I explore and defend the idea of privacy protections based on consumer protection law. In these sections, I do not draft legislative provisions; rather, I spell out a series of privacy principles that might underpin such laws. In the third section, I examine extra-legal measures, such as digital literacy education, market-based strategies and coding, which often seek to reinstate the power of user consent. Here, I further argue that privacy is not just about control and consent: other values and tests come into play, including respect. Finally, in the fourth section, I bring together these elements by arguing that the protection of privacy requires more than national or state-based approaches. The internet is international, *par excellence*, meaning that local solutions must work together with an international framework. Hence I touch on two further notions addressed by Kant: the realm of ends and cosmopolitanism.

The practical solutions I propose are derived from the descriptive and normative work done in previous chapters. They are also informed by interviews. This primary research comprises five interviews with: Timothy Pilgrim, Australia's privacy commissioner; Samantha Yorke, public policy and government relations counsel at Google Australia; Benjamin Carr, chief privacy officer at Australian telecommunications company Telstra; Nigel Waters, committee member of the Australian Privacy Foundation; and a spokesperson for the Federal Attorney-General's Department. These interviews, which were often extensive, are not presented in full. Rather, I employ key quotes to represent the main points made by each interviewee. (More details about the methodology and justification behind this primary research is included as an appendix.) As the identity of these interviewees suggests, this chapter has an Australian focus, but it is also outward-looking, containing discussion of laws and approaches in a variety of jurisdictions. It is beyond the scope of this thesis to be comprehensively global.

Ultimately, however, the aim is to point towards the outline of an international solution. If privacy is well protected, I argue, we will be able to move towards an internet that is richly cosmopolitan, and decidedly more utopian than dystopian.

I - Practical solutions: applying the theory

In the context of internet privacy, as we have seen, the notion of individual consent can be problematic. So much so that there are those who suggest that notice-and-consent should be abandoned as our go-to protection in favour of trust (e.g. Taddei and Contena, 2013), or harm minimisation (e.g. Wright and Raab, 2014), or user empowerment (e.g. Yorke, 2015). Certainly, these notions have merit. In a sense, trust is a complement of consent. Trust can only be established if a user's will and autonomy has been respected, and if that person has not been treated merely as an IP address ripe for harvest. Arguably, however, trust is harder to regulate than consent. Harm minimisation, by contrast, takes a consequentialist approach, which is a significant departure from consent. It does have the benefit, however, of considering the effects of privacy violations beyond the individual. It can extend the notion of privacy beyond mere data protection (as tends to dominate European privacy debates), to examine the risks or harms not only to the individual, but also to other individuals, to groups and to society as a whole (Wright and Raab, 2014). This fits my account of privacy as relational. Google, meanwhile, has adopted the phrase "user empowerment tools" as a mantra (Yorke, 2015). Evoking notions of autonomy and self-determination, the phrase is clearly consistent with consent. *Prima facie*, it would seem, a consenting user is an empowered user, and Google declares that consent remains its guiding principle, as we shall see. It is beyond the scope of this thesis to explore trust, harm minimisation or user empowerment in detail. Rather, I will concentrate on consent, as derived from Kant's formula of humanity, before broadening my focus in sections three and four.

i. Individual consent

The two-tier system of consent I have been advocating begins with individual consent, and hence with the notion of control. The principle is straightforward: if

a person (or persons), a company or a government intends to act in a way that significantly affects the privacy of an individual (see chapter five), there is an ethical obligation incumbent upon that person, company or government to seek that individual's consent.

It is no surprise, then, that individual consent, and the attendant notion of control, have been and continue to be highly significant when it comes to determining the limits of privacy online. As privacy advocate Nigel Waters says:

I strongly believe that as far as possible individuals should be given choices. In the internet context it's information privacy that we're talking about, and certainly maximising the opportunities for consent is desirable (Waters, 2015).

Australia's privacy commissioner, Timothy Pilgrim, agrees (at least for now):

There's been a debate going on for a number of years now about the value of consent, and more importantly for the value of notice, around the collection of information. There's one school of thought that there should be harm minimisation, so organisations should be able to collect information, and it should be handled in a way that recognises the potential harm to an individual from the use of different types of information. But I think until that debate has matured, there's still a very, very strong role for notice and consent in the collection and handling of personal information (Pilgrim, 2015).

Google too sees the merits of consent:

Google considers any information that a user shares with Google, whether by Gmail, Docs, Calendar, or privately shared content on Google+ or YouTube to be private. This includes search history too, whether a user is signed-in or not. We don't sell this information or share it with advertisers *without consent* (Yorke, 2015: italics mine).

Google's stated commitment to consent involves giving users as much control as possible. "Google believes that individual users are best placed to make decisions around how their data is managed" (Yorke, 2015). As such, specific tools had been created by November 2015, including: Dashboard, a single page on which users can see the products and services on which their information is stored; Account Settings, to let users manage which data they start and stop sharing with Google; and Ad Settings, which includes the capability of opting out of interest-based ads. These are supplemented by other general and specific privacy protecting mechanisms built upon the principle of control: "Being transparent and providing users with individual controls to manage their privacy settings is critically important at Google. We have made a significant investment in developing tools and controls for users" (Yorke, 2015). Of course, simply making consent available is not enough. If the option of consent is hard to find, or if

default settings are public rather than private, then arguably the possibility of consent, as defined in chapter five, has not really been offered (I return to default settings below). Google argues it has sufficiently built consent/control into its privacy policy in the shape of the My Account pages, launched in 2015 with the aim of giving users control by making privacy settings simpler and easier to use (Yorke, 2015).

When offering consent, companies such as Google must also consider the competence of users. In chapter five, I argued that whenever the issue of individual consent arises, the first issue involves competence to consent. If an individual is not competent, then hypothetical consent will generally be required, in the form of imputed consent. The internet, as we have seen, is remarkable for the durability and the searchability of its contents. If a child of 13 or a man in a manic phase of bipolar illness consents to the sharing of a potentially damaging image on social media, is that consent morally justifying? Presumably, no, because these people are not competent to consent, and hence no real possibility of consent has been offered. Generally, if someone is not competent to consent, we need to turn to consent by proxy, and look to a parent, or guardian, or carer to provide consent or dissent in that individual's stead. Otherwise, the 13-year-old has too great an ability to harm her future adult self, just as the mentally ill man is prone to endangering his future recovered self. Alternatively, in cases such as drunkenness, a proxy is not the answer. Rather, we must await the return of sobriety. More precise definitions of competence ought to be enacted via the collective consent of the law (as discussed in the next section). Detailed laws and policies are required to spell out not just what qualifies as consent, but also who is competent to consent, and what procedures must be followed in cases of incompetence.⁵⁹

⁵⁹ The privacy ethics that ought to prevail for 13-year-old social media users deserves a thesis of its own. Here, I have made the claim that 13-year-olds are, in many cases, not competent to consent. Clearly, this flies in the face of current norms. Thirteen-year-olds can and do share huge volumes of personal material on social media (see chapter two). At the very least, for a 13-year-old to be judged competent, a number of factors ought to exist. First, the child ought to be aware of the risks. Here, education can help. Second, settings ought to be private by default, public by choice, rather than *vice versa*. And third, consent ought to be iterative and layered, with frequent reminders about the nature and effects of sharing. These factors are explored below, as is the suggestion that data shared by minors should be "sealed", with one's data profile effectively wiped clean at age 18.

Once it has been established that a person is competent to consent, the best conception of consent involves actual consent, defined also to allow for possible consent. If we are to look for practical solutions, we must look for ways to enact and enforce this conception. In chapter five, I proposed three necessary conditions. First, consent must be *bona fide*, and to the best of the knowledge of the person or entity seeking consent, thus providing for a *caveat venditor*, rather than *caveat emptor*, approach. This aligns with Kant's emphasis on intention. Hence the focus is on Google, or Facebook, or the NSA, to offer consent to the best of their knowledge. Second, consent ought to take into the account the particularities of the person whose consent is sought. If social media is attracting a lot of teenage users, then consent requirements must be adjusted accordingly. Even if teenagers are competent, they are more *vulnerable* than adults. And third, consent must be iterative and layered. Each of these conditions ought to be codified in the collective consent of the law. One approach, as I discuss in section two, involves the provision of general protections, analogous to those found in consumer protection law. Alternatively, these conditions might be supported and enforced through more specific legal protections.

The third condition is particularly significant in an online context. Entities such as Google have already taken steps to make consent ongoing and responsive. This also involves giving users choices that are simple, clear and concise. Google's global privacy counsel, Peter Fleischer, says that a clear, logical flow is required for notice-and-consent provisions because that is how humans think of privacy, and also because more and more online interactions are taking place on mobile devices with relatively small screens. As Fleischer says:

The big innovation in privacy notices at Google is that they come up when that issue first arises. It's incremental. Mobile is the ultimate test – we need to get notices down to a sentence. A 10-page notice form is a catastrophe (Fleischer, 2015).

If you are using an app to help you find the nearest Mexican restaurant, your smartphone might ask you in one short, plain phrase if you agree to turning your location services on, because that particular app requires you to do so for it to function. This is user-friendly. What's more, the question, once answered, should not be considered settled forever. Even if a user agrees to turn on her location services, she needs to be reminded of that choice, and to be given the option to change her mind. What's more, technology is changing quickly; many platforms

are, it seems, in a state of perpetual reinvention. Instagram was launched in 2010 as an app for applying retro filters to photos taken on phones; in 2012, it was bought by Facebook for \$1 billion; since then, it has grown into a photo- and video-sharing site dubbed “the social media tool that defines the millennial generation” (Bruner, 2016). The notice-and-consent provisions agreed to in 2010 would have little relevance now. The issue of consent should arise when a company changes what it does, just as it should arise when a user is trying to do something new. For Nigel Waters too, consent is more likely to be meaningful if sought at the point of use, rather than in a larger, more general provision offered upfront (Waters, 2015). This aligns with research showing that users sometimes need to be nudged into caring about their privacy (Almuhimedi et al., 2015). It seems that any genuine attempt to offer the possibility of consent online requires a series of specific iterations.

ii. Collective consent

What about the law? As James Grimmelman asks, “Is the loss of privacy in social media something lawmakers ought to worry about and, if so, what should they do?” His answer? “A clear yes: users want privacy, deserve privacy, and cannot easily secure privacy for themselves” (Grimmelmman, 2009: 795). Grimmelmman’s focus is social media; but his point applies just as well to the internet more generally. It is from the law that the right to privacy derives. And online, as we have seen, the right to privacy is often unclear or weak. New regulatory strategies are required - particularly in countries such as Australia, where privacy is under-regulated - in order to provide an adequate response to potential harms (Meese, 2015: 144).

With the formula of humanity as our ethical foundation, individual consent is a logical starting point for the protection of privacy. However, individual consent is simply inadequate to protect privacy on the net (see chapter five). I have thus been arguing that this brings us to the Kantian notion of the “united will of the people”, or collective consent, in the form of just laws and regulations. It is a view supported by Waters: “There’s an awful lot of context in which it [individual consent] is simply not practical, and that’s why there needs to be some collective decision upfront about what is permissible and what isn’t” (Waters, 2015). This is

evident, for instance, with mobile phones and other locative media. The potential privacy impacts of location-based tracking are dramatic, and demand regulation: “We need to recognise that laws, policies and regulatory frameworks can help” (Dwyer, 2015b: 133). Often, it is the secondary uses of data that require regulation. For Amitai Etzioni, these secondary uses mark out the transition from the paper age to the digital age, a transition more dramatic than from the hand grenade to the nuclear bomb (Etzioni, 2015: 1263-1264). Given these secondary uses, internet law scholar Viktor Mayer-Schönberger agrees that informed consent isn’t enough. For Mayer-Schönberger, the concept of informed consent implies individual responsibility and self-determination; unfortunately, it can fail to protect the powerless (Mayer-Schönberger, 2015). In the context of online privacy, regulation is urgently required because there exists a strong power imbalance, there is complexity and there are externalities. These externalities include the situation where I agree to supply my DNA to a database. Yes, I have consented, but my family have not, even though, just by supplying my DNA, I have revealed a great deal about them too. Mayer-Schönberger’s prescription is regulation: “We need to regulate the use of personal data, and envision a framework for the responsible use of big data” (Mayer-Schönberger, 2015; see also Mayer-Schönberger and Cukier, 2013: 173-174). In line with the access model of privacy I have been advocating (see chapter three), individual consent comprises the user exercising control, and collective consent comprises externally-imposed restrictions on access. Like Mayer-Schönberger, I propose that collective consent in the form of regulation is urgently required for internet privacy.

Of course, privacy regulation already exists. In fact, it exists in abundance. However, it varies tremendously. Within countries, it varies from jurisdiction to jurisdiction. Between nations, it varies even more. In chapter three, I contrasted the legal protections for privacy in Europe (strict) with those in Australia and the United States (lax). In Australia, privacy regulation is interstitial and chaotic, a mix of federal, state and local law that is *ad hoc* and sometimes ill-fitting. Nowhere, in any of these instruments, is “privacy” defined. A first suggestion is that privacy would be better served by regulation that is more coherent and codified. I return to this point in section two. A second suggestion is that in all jurisdictions a logical legal starting point would be to define privacy (in terms of

access, as I have argued), and then to prescribe, in both general and specific legislative instruments, when consent must be obtained.

The most obvious function for collective consent is *to reinstate and reinforce individual consent*. More specifically, the law can clarify and buttress individual consent, by helping to delineate the circumstances in which consent is required, and by defining what constitutes legally valid consent. In certain instances, this is what the law does. In Australia, the *Privacy Act* has been drafted to protect “personal information”, but offers stricter protections for the subset of personal information that is “sensitive information”. Sensitive information is defined as information pertaining to matters such as sexual practices and preferences, religious and philosophical beliefs, membership of trade unions and associations, and health information. To collect sensitive information, an individual’s consent is required (OAIC, 2017c). In other words, Australia’s *Privacy Act* is an attempt to formulate a general protection for “sensitive information” by making consent mandatory. In the United States, a series of federal statutes address privacy, including the *Health Insurance Portability and Accountability Act* (HIPAA) of 1996, which governs medical records, and the *Family Educational Rights and Privacy Act* (FERPA) of 1974, which governs the privacy of students (CDT, 2008). These instruments tend to reinforce individual consent. FERPA, for instance, covers the release of information from student educational records, provides for student access to their records and establishes a means for students to seek amendment of records they believe are inaccurate, misleading, or otherwise in violation of their privacy rights. The general principle is that personally identifiable information regarding a student cannot be disclosed without his or her written consent (Young, 2015: 562-564).

European law protects privacy much more vigilantly. Generally, when it comes to privacy and data protection, Continental strictness contrasts Anglophone lassitude. This is evident in European legislative instruments that reinforce consent. In Europe, the General Data Protection Regulation, or GDPR, will bind all member states once it comes into effect on 25 May 2018 (OJEU, 2016). The GDPR imposes onerous requirements for consent, thereby shifting the burden of responsibility and accountability from users of data to “controllers” of data (Rotenberg and Jacobs, 2013: 632). Under Article 6, “... processing [of data] shall

be lawful only if and to the extent that ... the data subject has given consent ...” (OJEU, 2016). Under Article 7, written consent must be presented “... in an intelligible and easily accessible form, using clear and plain language” (OJEU, 2016: 37). The issue of competence is addressed, *inter alia*, in Article 8 (1), which prescribes that “... the processing of personal data of a child below the age of 13 years shall only be lawful if and to the extent that consent is given or authorised by the child's parent or custodian” (OJEU, 2016: 37). The GDPR is general and sweeping (see also chapter three). Meanwhile, more specific pieces of legislation have also sought to reinstate the power of consent. In recent years, jurisdictions including California, Japan and the UK have criminalised revenge porn. In the UK, a law came into effect in April 2015 making it illegal to disclose a “private sexual photograph or film” without the consent of the person depicted, and with the intent to cause them distress. The penalties include a maximum of two years in prison (Barrett, 2015). Previously, the scope existed for prosecutions under existing obscenity laws, and civil remedies existed too, but few prosecutions took place (Barrett, 2015). Other jurisdictions, including in Australia, have been debating whether to criminalise the nonconsensual sharing of explicit images and videos (Godfrey, 2013). Revenge porn legislation is a clear example of jurisdictions clarifying and buttressing individual consent in particular types of cases by criminalising any failure to respect it. Similar legislation might reinforce the role of consent in the collection of biometric information. As we have seen in chapter two, huge strides are being made in voice recognition, facial recognition and other biometric technologies (Shankar et al., 2016). Such technology has enormous privacy-compromising potential, as do wearables and embedded microchips. Their impacts must be carefully deliberated and appropriately regulated (Michael and Michael, 2013). Privacy, I suggest, requires a mix of general and specific protections.

Often, then, the law works to support individual consent. In other cases, however, the law can and does *set the limits* of individual consent. In this way, it can override individual consent and dissent. These are situations where the law has deemed that user control is irrelevant, and that there must be externally-imposed restrictions upon access. Such cases were discussed in chapter five in the form of the Naked Rambler and the Overheard Diners. The former wanted to waive his rights to privacy by being naked in public. He consented to his privacy being

encroached upon. However, UK courts have overruled his consent, finding that he is not entitled to ramble naked, and must maintain his privacy. In the latter (hypothetical) case, two diners did not consent to their conversation being shared. The content of their conversation, however, involving the discussion of a child pornography website, meant that any accidental eavesdropper would be legally obliged to disregard their dissent and report to police. The same is true online: regulations can override consent and dissent. There are laws limiting what personal images I may share with whom, just as there are laws mandating the reporting of certain digital communications and behaviour. Again, such regulations can be general or specific. Laws have an especially important role to play for the vulnerable, including children (see section two, below). Already, children have become the first major focus of privacy regulation. Laws protecting children's privacy include the *Children's Online Privacy Protection Act* [COPPA] in the US, which includes anti-tracking provisions. Enforcement agencies include the US Federal Trade Commission. In Australia, advocacy groups include the Australian Communications and Consumer Action Network [ACCAN] (Dwyer, 2015b: 123).

One way in which the law can *both enforce and limit* the role of individual consent is by implementing a tort under which users can sue for breaches of privacy. There are civil causes of action for serious invasions of privacy in the US, the UK, Canada and New Zealand, created sometimes via statute, sometimes via common law (ALRC, 2014). In 1960, William Prosser argued that there were four distinct privacy torts, under which individuals can sue on four grounds: intrusion of solitude, which can involve physical or electronic intrusion; public disclosure of private facts, which involves sharing truthful private information in a manner which a reasonable person would find objectionable; the publication of facts which place a person in a false light; and the appropriation of a person's name or likeness to obtain benefit (Prosser, 1960: 389). In response, Edward Bloustein argued in 1964 that tort cases involving privacy involve a single tort (Bloustein, 1964: 1000). The cases he cites include *Pavesich* (122 Ga. 190, 50 S.E. 68 (1905)), in which a man sued successfully for an invasion of privacy after his image was used in an advertisement without his permission (Bloustein, 1964: 986). As Judge Cobb wrote:

Liberty includes the right to live as one will, so long as that will does not interfere with the rights of another or of the public. One may wish to live a life of seclusion; another may desire to live a life of publicity; still another may wish to live a life of privacy as to certain matters and of publicity as to others ... (quoted in Bloustein, 1964: 1002)

In this formulation from 1905, there is an invocation of control (“the right to live as one will”), tempered with reference to the rights of others and the public. It is an approach compatible with the prescription I have given of individual consent and collective consent, and of privacy as relational. In modern US tort law, Prosser’s four distinctions still prevail, arguably at the expense of the law’s ability to adapt to new technologies, given its lack of a single guiding, unifying concept (Richards and Solove, 2010). Nonetheless, there have been famous cases. In 2016, Terry Bollea, aka wrestler Hulk Hogan, was awarded US\$140million after suing the Gawker website for posting footage of him having sex with a friend’s wife. After Gawker was shut down, Bollea settled for \$31million. The website’s fans declared a dark day for press freedom; Bollea’s supporters celebrated a victory for privacy (Ember, 2016).⁶⁰ Any such tort, I suggest, must be limited by a public interest test. In the Bollea case, there appeared to be no public interest, only prurient interest, in publication. Had Bollea been shown committing a crime, however, the public interest might have trumped Bollea’s privacy rights. (The public interest is further discussed below.)

In the UK, a tort of privacy has crystallised only recently. It first emerged following the enactment of the *Human Rights Act 1998*, which incorporated provisions of the European Convention on Human Rights into domestic law. The application of these provisions led courts to extend the equitable action of breach of confidence, with some disagreement as to whether the action had thus in fact coalesced into a tort (Mo, 2017: 90). Then, in the 2015 case of *Google Inc v Judith Vidal-Hall* ([2015] E.W.C.A. Civ 311), three parties successfully sued Google for the misuse of private information by its use of internet cookies, given that the cookies were installed without consent. This was contrary to Google’s claims that user-generated content could not be tracked without the user’s permission (Mo, 2017: 89). The effect of *Vidal-Hall* is dramatic: the misuse of

⁶⁰ I am recommending a tort of invasion of privacy where none exists, but a tort is certainly no silver bullet. The Bollea case was bankrolled by Peter Thiel, a Silicon Valley billionaire who had earlier been outed as gay by a Gawker blog (Ember, 2016). Civil suits are notoriously expensive, thereby favouring rich claimants. As such, a tort must be supplemented by other remedies.

private information is now explicitly recognised as a tort in English law (Mo, 2017: 88-89).

Australia has no analogous tort, although federal and state reports have consistently recommended one, including in 2014 (ALRC, 2014: at 13 for a discussion of previous reports) and 2016 (NSWSCLJ, 2016). In its 2014 report, *Serious Invasions of Privacy in the Digital Era*, the Australian Law Reform Commission [ALRC] recommended a statutory civil cause of action, in which consent and the public interest play key roles. The tort would be:

... directed at invasions of privacy that are serious, committed intentionally or recklessly, and that cannot be justified in the public interest. It is also confined to invasions of privacy either by intrusion upon seclusion or by misuse of private information ... (ALRC, 2014: 6)

These elements of the cause of action (serious, intentional or reckless, not in the public interest, committed either by intrusion or misuse) must all be satisfied. Further, damages may be awarded for emotional distress. The ALRC also recommended that consent feature as a defence. In other words, anyone being sued for a serious invasion of privacy would be able to mount a defence that consent had been obtained (ALRC, 2014: 8). Other defences include lawful authority, necessity and fair reporting. (ALRC, 2014: 7-8). Crucially, the right to privacy protected under the proposed tort would be limited by a robust conception of the public interest. As the ALRC recommended: a court must be satisfied that “the public interest in privacy outweighs any countervailing public interests” (ALRC, 2014: 7). The ALRC thus proposes a law requiring:

... a crucial ‘balancing exercise’, in which courts weigh privacy against other important public interests, such as freedom of speech, freedom of the media, public health and safety, and national security ... A plaintiff should not be able to claim that a wrong has been committed - that their privacy has been seriously invaded - where there are strong public interest grounds justifying the invasion of privacy (ALRC, 2014: 7).

In this way, the ALRC’s recommended tort seeks to balance individual control, including that conferred by individual consent, with externally-imposed restrictions upon access, comprising the collective consent as determined by legislators and judges. Such a tort of privacy would have the potential to provide significant general protections for privacy, allowing for courts to adapt to significant advances in technology. Indeed, the ALRC based its decision to recommend a tort on nine principles, which included that privacy laws: should be

adaptable to technological change; should be clear and certain, coherent and consistent; and should make justice accessible to all. Recognising that privacy is a fundamental value worthy of protection, the ALRC noted also that privacy protection is an issue of shared responsibility between individuals, industry and government (ALRC, 2014: section 2). These are compelling arguments.

In response to the ALRC, Google recommended that a tort should only be available to: natural persons; where a person has a reasonable expectation of privacy; where the act is sufficient to cause substantial offence to a person of ordinary sensibilities; and where the act is intentional or reckless. Moreover, Google agreed that a defence of consent be included in the tort (Meese, 2015: 143). A tort of invasion of privacy, enacted by parliament and refined by the courts, can go far to protect privacy. In line with my arguments below, the tort could specify different criteria and remedies in the face of threats from individuals, companies and governments. After all, not all encroachments are equal. A vast difference exists between a hacker gaining unauthorised access to a webcam for personal gain and an AI system employed by the government for national security that automatically records webcam footage no human will ever see, unless a court orders otherwise. To duly protect privacy, there would still be the requirement for other laws, including criminal laws. A tort is merely one potential ingredient among several. However, a carefully-worded tort of privacy would align well with the normative model articulated in chapter five, in which, following the formula of humanity, individual consent is both reinforced and limited by collective consent.

iii. The threat from individuals

In chapter two, I identified a triple threat to privacy: from individuals; from organisations; and from governments. In many ways, these threats overlap; but in some ways, they are distinct. For instance: an individual, unlike a company or government, may face criminal penalties, including time in jail, for breaches of privacy; a company, unlike an individual or a government, might effectively be prompted to respect privacy via market mechanisms; and the government, unlike a company or an individual, may be justified in compromising an individual's privacy in the interests of national security. Some current laws recognise this

distinction. Australia's chief privacy-protecting law, the federal *Privacy Act 1988*, only applies to relevant "entities", defined as "most Australian and Norfolk Island Government agencies, all private sector and not-for-profit organisations with an annual turnover of more than \$3 million, all private health service providers and some small businesses ..." (OAIC, 2017a). As such, the Act does *not* cover individuals, meaning there is no onus under the *Privacy Act* on individuals to obtain consent when collecting sensitive information. The gap in Australia's *Privacy Act* is emblematic of regulatory blind spots, which are common internationally (although much less in Continental Europe, as we have seen). These are particularly common, it seems, when the threat comes from individuals.

As we have seen, one example is "revenge porn", for which authorities have been reluctant to prosecute using outdated, ill-fitting provisions (Gatford, 2015). In Australia, these include the federal crime of "using a carriage service to menace, harass or cause offence" (ALII, 2017b: division 474). Revenge porn reveals how individual consent and collective consent ought work together, and also how a tort of invasion of privacy can help. In 2011, following a relationship breakup, Danish student Emma Holten found that nude photos of her had been maliciously posted on the internet.⁶¹ She then received hundreds of messages from men around the world, many of them abusive.

Suddenly, I noticed that this dynamic – sexualisation *against* her will – was everywhere. Take 'creepshots', a global phenomenon which entails photographing women without their knowledge or consent, in order to share them in a sexual context online ... Here, again, women are used as objects whose lack of consent, of participation, provides the reason and allure of their sexualisation (Holten, 2015).

Two years after the photos were posted, Holten commissioned a photographer to take another series of nude photos, which she then posted online. It was an attempt to make her a sexual subject instead of an object. "Consent is key. I did this. Just as rape and sex have nothing to do with each other, pictures shared with and without consent are completely different things" (Holten et al., 2015). Such "image-based abuse" is common: an extensive 2017 study found one in five Australians have been victims (Henry et al., 2017). Holten's case illustrates how she (and others) have sought to reclaim consent/control in a digital environment

⁶¹ In Holten's case, it may have been a hacker, not an ex-partner, who shared the images non-consensually. Hence her case, strictly speaking, may not be "revenge porn". However, the analysis still holds, and it has been suggested that a better phrase is "image-based abuse" (Henry et al., 2017)

where it has been taken from them. Certainly, issues of coercion arise, given that Holten was sharing nude photos only after other nude photos had been shared non-consensually. What's more, given the extent of the problem, Holten's case also reveals that collective consent has a significant role to play. In various jurisdictions, revenge porn has been criminalised (Matsui, 2015; Franks, 2015; Barrett, 2015). In this way, the law can punish nonconsensual sharing of intimate material. As a supplement, a tort can provide a civil remedy and pecuniary damages (Gatford, 2015).

In line with the formula of humanity, an effective regulatory approach to protecting against the threat from individuals, including the threat of revenge porn, would involve a blend of specific and general laws that reinforce and/or limit the role of individual consent. A general law might include a tort of privacy, implemented ideally through the legislature rather than case law (ALRC, 2014). General laws might also consist of consumer-style protections, as discussed in section two. Indeed, a tort of privacy would do well to *encompass* consumer-style protections. Meanwhile, specific laws could target specific abuses, including revenge porn and online impersonation. Whether general or specific, laws are also needed, it seems, against the threat of technology such as drones. In 2017, an Australian woman swimming naked in her backyard pool was distressed to see a drone above her. "My fences are really high and secure and there's big trees around the backyard, so it's the last place you'd think your privacy would be able to be invaded," she said. In response, a legal scholar argues new laws are urgently needed, as it is generally lawful in Australia for a person to film what they can see from a public space, including from the sky (Mitchelson, 2017a).

The threat from hackers, outlined in chapter two, is substantial. Suicides followed the Ashley Madison hack of 2015, which exposed the identities of users of the adultery website (Segall, 2015; Vallor, 2016: 192). By accessing private data, hackers might steal money, blackmail users or harm people in more subtle ways. One team of researchers cited the case of "Brian", whose Facebook profile was hacked multiple times. The first time, the hacker altered Brian's "interested in" to insinuate that Brian was gay. Later, the hacker changed Brian's relationship status to "I'm having a hard time coming out of the closet right now." Brian deleted his profile and quit Facebook, only to be hacked again after returning months later

with a new email address (Debatin et al., 2009).⁶² How does the law deal with such a case? How should the law deal with such a case? Clearly Brian did not consent to this invasion; clearly he was treated merely as a means. In many jurisdictions, such hacking is a criminal offence. Had these events occurred in Australia, the hacker could have been prosecuted under division 477.1 of the federal *Criminal Code Act*, which prescribes that a person is guilty if she makes “any unauthorised access to data held in a computer; or any unauthorised modification of data held in a computer; or any unauthorised impairment of electronic communication to or from a computer ...” (ALII, 2017b). In 2016, following the celebrity nude hack of 2014, a Pennsylvania man was sentenced to 18 months in jail after pleading guilty to one count of unauthorised access to a computer to obtain information under California’s *Computer Fraud and Abuse Act* (Yugas, 2016). Similarly convicted of unauthorised access (and extortion), Cassidy Wolf’s hacker was also sentenced to 18 months in prison (Gander, 2014). When it comes to the invasion of privacy by individuals, including hackers, the criminal law has a role to play, but criminal sanctions can be effectively supplemented with options for civil redress, including a tort of invasion of privacy (Calkins, 2000: 223).

iv. The threat from companies

In many ways, the threat from individuals pales beside the threat from companies and governments, who have a far greater capacity to collect, store and sort information. There is thus a significant role for regulation which specifically targets companies, as there is for regulation which specifically targets governments. For the threat from companies and governments, a tort could play an especially significant role.

One recent development specifically targeting breaches by companies is Europe’s new “right to be forgotten”, or “right to erasure”, which allows for individuals to request that companies such as Google remove specified links to URLs from its search results. It’s a striking development that reveals our times: “The ‘right to be

⁶² To complicate matters, this is no typical privacy case. Brian was not, in fact, gay. It was his *invented* privacies that were exposed. Any law protecting privacy would need to specify whether it covers invented privacies, or whether such cases ought to be covered by more general protections against, say, unauthorised access and online impersonation.

forgotten’ is in some senses the 21st century version of the ‘right to be let alone’” (Dwyer, 2015a: 184). The new right derives from the judgment in a 2014 European Court of Justice case, *Google Spain SL & Google Inc v Agencia Espanola de Proteccion de Datos (AEPD) & Mario Costeja Gonzalez* (Case C-131/12, 13 May 2014; see discussion in Kranenborg, 2015). A Spanish man, Mario Gonzalez, objected to the fact that whenever someone Googled his name, the search results would include links to two newspaper stories from 1998, describing real estate auctions prompted by proceedings to recover his social security debts. In 2010, Gonzalez lodged a complaint with the Spanish data protection authority against the newspaper and against Google, requesting that these links be removed from search results, citing the European Parliament’s Data Protection Directive 95/46/EC on the protection of individuals with regard to processing of personal data and on the free movement of such data. After the Spanish courts had considered the case, the European Court of Justice found in Gonzalez’s favour, acknowledging that a balance needed to be struck: between the right of the user demanding privacy, and the right of internet users to information. The court noted:

That balance may ... depend, in specific cases, on the nature of the information in question and its sensitivity for the data subject’s private life and on the interest of the public in having that information, an interest which may vary, in particular, according to the role played by the data subject in public life (quoted in Kranenborg, 2015: 73).

The court also noted that accurate data that is lawful initially may over time become incompatible with Directive 95/46. In such cases, the links should be erased. The court thus found that Gonzalez did have, in this case, a right to be forgotten, and that accordingly Google should remove the links breaching the provisions of the directive. As Google’s Peter Fleischer says, “We were defending the principle that as long as content is legal on the web, you should be able to find it through a search index. The court disagreed” (Fleischer, 2015).

In the year following the Gonzalez decision of May 2014, Google received more than a million requests to remove links, and approved roughly 40 per cent (Fleischer, 2015; Calpito, 2015). It is important to reiterate that where the request is approved, the websites are not removed from the web. Rather, the relevant links do not appear in search results. However, Gonzales would need to approach other search engines if he wanted them to remove those same search results. Nonetheless, European courts have sought to enforce the right. In March 2016, the

French data protection authority fined Google €100,000 for not scrubbing web search results widely enough. Rejecting Google's arguments, the court found that the right to be forgotten entitled users to have links removed from the international "google.com", and not just France's "google.fr" and Germany's "google.de" (AAP, 2016b). Europe's recognition of a right to be forgotten clearly reinstates a degree of user control. Further, it also recognises, as noted above, that this right must be balanced against "the interest of the public in having that information", and that this interest "may vary ... according to the role played by the data subject in public life" (Kranenborg, 2015: 73). A person in public life, it would seem, has less of a right to be forgotten. Following the Gonzalez decision, the right to be forgotten, more properly known as the "right to erasure", will be enacted in the GDPR, which comes into effect in May 2018 and replaces Directive 95/46/EC. Article 17 of the GDPR entitles the data subject to have the data controller erase his or her personal data, cease further dissemination of the data and even have third parties halt processing of the data. This may happen when the data is no longer suited to original purposes for processing, or if the data subject withdraws consent. Again, this must be balanced against the "right of freedom of expression and information", and also for reasons of public interest in public health and archiving purposes (OJEU, 2016: 43-44). Hence the GDPR, as an act of the European Parliament, is an act of collective consent that seeks to re-empower individual consent, but also to set the limits of individual consent when weighed against other rights and considerations.

It has been argued that the right to be forgotten is useful in the search engine age, and indeed chimes well with Nissenbaum's theory of contextual integrity; however, it has been suggested that the binary approach of "forgetting" versus "remembering" should be improved by allowing for a more nuanced approach which also allows for "delisting" and "reordering" (de Mars and O'Callaghan, 2016). A more nuanced approach, I suggest, would complement the nature of privacy on the internet, which is marked by subtleties and layering, including as a result of the multiplication of place. On the net, the right to be forgotten (or delisted, or re-ordered) is a promising response, according with the formula of humanity's command that consent is pivotal when the issue is whether we are treating someone merely as a means. The right to be forgotten is one potent

strategy against the threat from companies. It re-empowers individual consent, but also limits individual consent by balancing it against other rights and interests.

v. The threat from governments

I have argued that collective consent, in the shape of the law, can limit individual consent, and that sometimes this limit is in the form of the public interest. But is the law not in the public interest, *ipso facto*? Is not the government, as the representative of the people, enacting the public interest with each piece of legislation, just as each judicial decision and each executive order is also an expression of the united will? In this section, I consider the justness of laws and policies underpinning surveillance by government agencies, including with reference to possible collective consent, but also with reference to Kant's principles of publicity. This discussion continues below: in section two, where I consider how government agencies ought best to balance the right to privacy against other rights; and in section four, where I return specifically to the *international* nature of surveillance undertaken by government agencies. With this analysis, I aim to provide some normative guidelines for surveillance, while also proposing legal remedies based on the formula of humanity for citizens unjustly surveilled.

The revelations of Edward Snowden are described in chapter two. In 2013, as a contractor to the National Security Agency [NSA], Snowden leaked classified documents revealing the extent to which the intelligence agencies of the US, the UK, Canada, Australia and New Zealand – the so-called “Five Eyes” – were putting citizens under surveillance in the interests of national security. As Snowden told journalists:

The NSA specifically targets the communications of everyone. It ingests them by default. It collects them in its system, and it filters them, and it analyzes them, and it measures them, and it stores them for periods of time, simply because that's the easiest, most efficient and most valuable way to achieve these ends. So while they may be intending to target someone associated with a foreign government or someone that they suspect of terrorism, they're collecting your communications to do so. Any analyst at any time can target anyone ... (*Citizenfour*, 2014)

Much of this collection involves the internet. For instance, Snowden revealed that the UK intelligence agency GCHQ uses sophisticated tools to track people by impersonating spammers and monitoring social media postings, and that the

SPRING BISHOP program was created to find “private photos of targets on Facebook” (Ball, 2014). The first point is that if the government is acting illegally, then its actions are seemingly not in accordance with the united will of the people. If, for instance, the NSA was acting in contravention of the US Constitution, or any other law, it is unlikely to be enforcing collective consent. For now, the lawfulness of the NSA’s activities remains contested. “Is this legal? The real answer is that we don’t know” (Schneier, 2015: 65-67). In one 2015 case, the US Court of Appeals found that the bulk collection of telephone metadata was illegal (Roberts and Ackerman, 2015). Certainly, any government’s use of invasive spyware such as NSO Group’s Pegasus on its citizens would appear to be unethical; in many countries, it is clearly illegal (Perlroth, 2017: see chapter two). If the NSA, GCHQ, the ASD and analogous agencies are acting illegally, they are, *prima facie*, acting unethically.⁶³ In some cases, it seems, the NSA was acting entirely *in the absence of* relevant laws. This too is problematic.

For the sake of argument, let us assume the NSA and analogous government agencies are acting in accordance with the law. In this case, the question becomes: is the law under which they are operating just? As described above and in chapter five, the law seems to be, by its very enactment, an expression of the united will of the people. However, in some cases, the legislature may, whether erroneously or cynically, pass laws that are unjust. For Kant, the ultimate test of political legitimacy is a social contract, and the social contract is built on possible consent of its citizens. As Kant wrote, the social contract takes the form: “If a law is so framed that a whole people could not possibly give it their consent ... the law is unjust ...” (Kant, 1996b). If, however, it is possible that people might agree to it, then the people have a duty to regard that law as just, even if, on their current way of thinking, they would probably refuse to agree (Kant, 1996b). Not all laws embody the united will; only *just* laws do so; and a just law is a law to which all the people in that jurisdiction could possibly have assented. A law implementing slavery could never be a just law, as this could not be a law to which all citizens could possibly assent. Slavery is an easy example; so too is Kant’s example of legislated hereditary privilege (Rauscher, 2016). With privacy, however, the answer is not so readily apparent. My right to privacy is considerably more subtle

⁶³ Of course, if the NSA is breaching an *unjust* law, then it may not be acting unethically, just as a citizen is not acting unethically by breaching an unjust law.

and conditional than my right not to be enslaved or disadvantaged by lack of hereditary privilege. Some laws under which the NSA has been operating may well be unjust under the possible consent test, as I shall argue below, particularly once we take into account the international nature of such surveillance.

Apart from the “united will of the people” test involving possible consent, there is another test to determine which laws and policies are unjust: and that is the publicity principle. For Kant, government policies ought not be incompatible with publicity. That is, all actions relating to the rights of others are wrong if their maxim is incompatible with publicity (Kant, 1996c: 8:381). As Allen Wood writes:

... a political maxim of policy is known to be unjust when it is possible for the politician to foresee, before implementing it, that its being made public would arouse such public opposition that the aims of the policy would be defeated (Wood, 2014: 77).

This is, Wood argues, Kant’s first principle of publicity. This first publicity principle does not settle once and for all whether a policy accords with right or is contrary to right, but it does give us a sufficient condition. If a policy breaches this principle, it cannot be just. As Wood writes, the concern here is “a political maxim of policy”: it is general policy, not the program specifics, that must be capable of being made public. In the case of the NSA, this maxim might be: we reserve the right to inspect anyone’s communications and data if we deem it relevant to national security, and this includes the right to store and record everyone’s communications and data, because we don’t know what might be relevant in the future.

The NSA has been working largely in secret, under executive orders, laws and court decisions that were sometimes covert (Schneier, 2015: 65-66). An important distinction must be drawn here between the legislature and the executive. On the one hand, all laws passed by the legislature ought to be public. On the other hand, the executive can, in certain cases, justifiably hold discussions and formulate plans in secret. In this way, a government can debate and devise policies behind closed doors. Even for executive secrecy, however, there are limits. Documents released by Snowden suggest that some of the policies under which the NSA and analogous agencies have been operating breached Kant’s first principle of publicity. Certainly, there was an absence of transparency in the way the NSA

was operating, with commentators arguing that greater transparency is required in the review process for classified programs (Schmidt and Cohen, 2013: 266). The policies underpinning the NSA's programs of domestic surveillance were not aired and debated before implementation. US citizens were unable to voice their opinions. If US citizens had been made aware of government surveillance and its security benefits, they may well have shown an overwhelming support for the underlying policies. This is worth considering (but would not, of course, settle the matter of whether these policies are just). As it happened, once programs and policies were revealed by Snowden, public outcry led to the passage of the *USA FREEDOM Act* in 2015, which sought to scale back the NSA's domestic surveillance capabilities. This response suggests that these policies were incompatible with publicity and hence unjust.

Transparency is a key ingredient of just laws and policies. We might even say that transparency, ironically, is a fundamental principle for privacy laws. Arguably, there is an even greater need for privacy laws and policies to be public than other types of laws and policies. The Australian *Privacy Act*, which protects against invasions by government agencies and big companies, makes transparency a key goal when it prescribes in Privacy Principle 1: "The object of this principle is to ensure that APP entities manage personal information in an open and transparent way" (OAIC, 2017b). As far as possible, openness is required about what is being collected, who is collecting it, for what purpose it is being collected, and (crucially) what secondary uses might follow. Elgesem writes that it is imperative that privacy principles be made public: "Privacy norms are part of 'the basis of social cooperation,' to borrow a term of Rawls; hence they are public principles" (Elgesem, 1996: 47). This is not to say that the commitment to transparency ought to be absolute. Governments and public institutions, like individuals, need to know that not every detail of their operations will be made public (Vallor, 2016: 205). The executive can, as noted above, have secrets. Generally, though, the principle of transparency, to which I return in the next section, is a crucial ingredient of privacy law and policy. Publicity enables citizens to reflect on whether a law enacts their will, and hence whether or not it is just. As Whitman writes, "Law will not work as law unless it seems to people to embody the basic commitments of their society" (Whitman, 2004: 1220).

It is logical that a government might want to put its populace under surveillance. The reasons may be self-serving (to help it stay in power), altruistic (to keep people safe), or both. Further, the government may reason that people will share more information about themselves if they do not realise they are under surveillance. Hence a conflict of interest arises: on the one hand, it is the government's responsibility to legislate to protect citizens' privacy; on the other hand, the government may benefit from invasions of citizens' privacy. Added to this is the way government surveillance can benefit from corporate surveillance. *Inter alia*, the government's role is to protect people's privacy from companies. However, if a government benefits from gaining personal data from companies, how can it be expected to regulate those companies dispassionately? Among Snowden's revelations was the PRISM program, under which the NSA collected citizens' internet communications from at least nine companies, including Google, Apple and Facebook. PRISM is at the heart of the NSA's data gathering programs (Greenwald and MacAskill, 2013: see chapter two). For governments, the allure of company-obtained data is self-evident. For Snowden, this has created a power imbalance between government and individuals: "We are building the biggest weapon for oppression in the history of mankind" (*Citizenfour*, 2014). At the least, we can see how a government might be tempted to invade the online privacy of its citizens. One obvious safeguard is for greater transparency. That way, the populace would have a greater ability to express its consent and dissent to laws and policies that affect individual privacy.

As I have already signaled, a further issue vexes the question of whether laws implementing government surveillance are just laws: the issue of an international mandate. The internet is international; so too is government surveillance. In this global surveillance network, an agency such as the NSA is, effectively, spying on the world's citizens (Schneier, 2015: 75-77; see chapter two). But how does someone in Australia possibly consent to the NSA's surveillance? In the digital age, does a just law require the possible consent of *everyone in the world*? The answer would seem to be yes. Clearly, ethics transcend national borders, just as the internet does. I return to this issue in section four. What's more, it must be stressed that the united will of the people test merely sets a minimum standard. It tells us which laws are *unjust*, and should thus be abandoned. However, it does not tell us which laws are *good*. It does not tell us which laws we *should* adopt.

Whether the laws and policies underpinning surveillance by government agencies such as the NSA are good laws and policies is a question beyond the scope of this thesis (though I touch on it in section two). My point is that we must consider whether NSA surveillance is based on just laws and policies, but we also need to consider, even if those laws and policies are just, whether a government *ought* to adopt them.

Finally, we must note that the challenges to privacy from governments have certain unique characteristics, which must be accounted for when devising forms of relief. As Bloustein wrote: “the forms of relief available against a government officer are to be distinguished from those available against intrusions by a private person” (Bloustein, 1964: 975). So then, how can privacy be protected against challenges from governments? First, a tort of invasion of privacy, built around the principle of consent, would allow for suits to be brought against governments, as well as against individuals and companies. And second, a requirement for government transparency can be enshrined in law, with limited but explicit exceptions. Such a requirement for transparency could effectively be built into a system of legal protections based on consumer law, to which I now turn.

II - Legal protections modeled on consumer law

Privacy is a coherent value, writes Ruth Gavison, but there is a lack of coherence in US judicial decisions about privacy (Gavison, 1980: 459-461). This, as I have shown, is true across various jurisdictions. What’s more, in many jurisdictions there is a corresponding lack of coherence in legislative approaches to privacy. Gavison argues: “There is much to be said for making an explicit legal commitment to privacy. Such a commitment would affirm that privacy is not just a convenient label, but a central value” (Gavison, 1980: 467). This contention underpins my argument. Privacy matters, but without legal protections its worth is easy to dismiss. Clear and coherent legal protections of privacy are a legal commitment, and also a powerful statement of its value. Above, drawing on my analysis of Kant’s ethics and politics, I have argued that privacy needs to be protected with a range of legal mechanisms, both civil and criminal. These mechanisms must take into account the distinct nature of threats to privacy from

individuals, companies and governments. What's more, these mechanisms can be formulated to align with Kant's prescription for consent, drawn from the formula of humanity's prohibition on using persons merely as means. More specifically, I have then proposed that these mechanisms ought to include a tort for invasion of privacy, worded so as best to accommodate the triple threat facing internet users.

To enact this range of civil and criminal remedies, the legislature has a key role to play. In Australia, as we have seen, the Australian Law Reform Commission has recommended that a tort of invasion of privacy be enacted by statute, rather than by judicial decisions (ALRC, 2014: 11-12). As the ALRC notes, statutory reform is more certain, proactive and faster than the interstitial changes wrought at common law (ALRC, 2014: 12). Even so, parliament is often slow. Reform of Australia's *Privacy Act* took nine years. "It's a long process," says the privacy commissioner (Pilgrim, 2015). This is a particular concern in light of the radical, breakneck evolution of the internet. In the nine years it took to reform the *Privacy Act*, Facebook grew from a teething one-year-old to a global forum with a billion-plus users. In those nine years, revenge porn and cyberstalking entered the vernacular. As Benjamin Carr says: "There will always be a role for regulation, but how do you regulate when it's moving so quickly?" (Carr, 2015) One response to the law's glacial pace is to include *general* protections, designed to accommodate new advances in technology. For Australia's privacy commissioner, general prescriptions are essential to provide the requisite flexibility:

I think that's the only way you can achieve that balance for those two competing things – free flow of information and [individual] control – also for allowing a law to actually not have to be changed too regularly, because it's almost impossible to change laws regularly (Pilgrim, 2015).

Above, I proposed a tort of privacy, which would go some way towards providing both a general and responsive protection for privacy.

Another option (which I propose as a complement, rather than an alternative) is that privacy law takes its cue from the consumer protection law of countries such as the US and Australia. This is the approach advocated by legal scholar James Grimmelmann:

... some of the lessons the law has learned in dealing with product safety could usefully be applied to the analogous problem of privacy safety. Unlike database regulations, which tend to focus only on the flow of information in itself, a product-safety approach can also

consider how people use social media ... [Hence we can] map the products liability doctrine onto the problem of making social media safe for privacy (Grimmelmann, 2009: 813).

Grimmelmann draws a parallel between physically safe products and privacy-safe social software. Privacy advocate Nigel Waters also believes in general prescriptions founded on informed consent:

The important point is that we shouldn't always try and reinvent the wheel. There are a lot of parallels in other areas, whether it's environmental regulation or consumer protection more generally, which privacy can draw on, and contribute to as well. It's been a constant frustration to me throughout my privacy career that there hasn't been better cooperation between privacy regulators and consumer protection regulators, particularly in terms of the principle and the attempt to provide *informed consent*. At the very least consumers should know what organisations are collecting about them and doing with that information and therefore be in a position to exercise some degree of influence or to kick back if they don't like it (Waters, 2015: italics mine).

In other words, Waters argues, the collective consent of the law ought to put internet users in a position where they are better able to give or refuse informed consent, and where that informed consent then has the law's backing. As Waters says, on the internet people tend to be overwhelmed by the data before them, including in privacy policies. "That's why there needs to be some collective decision upfront about what is permissible and what isn't" (Waters, 2015).

For Waters, privacy protections ought to comprise an interplay of individual consent and collective consent. Viktor Mayer-Schönberger, as noted above, also argues that regulation is urgently required in the digital age. We need, in short, to envision a framework for the responsible use of big data, and this involves introducing ex-ante protections, rather than the ex-post protections that individuals currently try to invoke only once there is a problem.

We should give up consent as our primary go-to mechanism, its almost monopolistic power should be replaced with regulation, we've already done this with seat belts, food safety, drug safety (Mayer-Schönberger, 2015).

Hence Mayer-Schönberger envisages a prescriptive and preventative model rather than a reactive and punitive model. To enforce it, a regulator is required, he says, just as a regulator exists for food, drugs and car safety. This model is potentially analogous to consumer protection law. For Grimmelmann, Waters and Mayer-Schönberger, general protections have a key role to play for internet privacy.

As noted in chapter three and above, privacy law in Europe, the US, Australia and elsewhere already incorporates a mix of civil and criminal remedies. These laws are sometimes specific, sometimes general, and enable different jurisdictions to value privacy to different degrees. The most stringent protections of privacy, as we have seen, are enforced in Europe, where the GDPR comes into effect in 2018, and includes the right to erasure. In Australia and the US, by contrast, privacy protections are *ad hoc* and significantly more limited. My claim is that privacy law *requires* a mix of specific and general, civil and criminal. Revenge porn, for example, seems to require specific remedies that ought to be supplemented by general redress, including the avenue of a tort for serious invasions of privacy. Available against intrusions by individuals, companies and governments, general protections have the benefit of responding nimbly to advances in technology. What's more, general protections can also serve as an effective deterrent. One strong example of a codified series of general protections is expressed in the GDPR. Another promising approach for general laws, which I now detail, is to mimic consumer protections in a way that encapsulates the formula of humanity's prohibition on treating others merely as means.

My project has been to apply Kant's formula of humanity to internet privacy. This involves the judicious application of individual consent. Even if we acknowledge it is no longer our primary go-to mechanism, as Mayer-Schönberger contends, consent still has a major role to play, particularly if we are devising privacy protections modeled on consumer protections. The consent principle underpinning such provisions might be expressed: **"Any acts that significantly impact upon a person's privacy must, where possible and appropriate, seek to obtain and abide by that person's consent."** Protections that enact such a principle would explicitly reinforce the control of the user. In its defence of consent, such a principle would help to vouchsafe the dignity and autonomy of the user. Clearly, such a principle would apply for Emma Holten, whose intimate photos were posted online (see above). Image-based abuse constitutes a clearcut breach of our first principle. The consent principle would also capture some of the practices regarding the hidden transfer of data to third parties. Even more clearly, such a principle would invalidate Facebook's practice of creating shadow profiles, which involve compiling dossiers of information on people who are not Facebook users (see chapter two). Similarly, it would invalidate the now-abandoned Google Buzz

(Grimmelmann, 2009: 823-826). For such users, there has been no actual consent, and no offer of the possibility of consent. What's more, the law ought make explicit that a preliminary issue is *competence* to consent. Once competence has been established, consent must satisfy three conditions. It must be to the best of the knowledge of the person or entity seeking consent; it ought to take into the account the particularities of the person whose consent is sought; and it ought to be iterative and layered, as required.

Beyond consent, there remains scope for the application of further principles. These include transparency and fairness, which both clearly link with the formula of humanity, and particularly its *prima facie* prohibition on deception. Above, I have discussed the link between transparency and just laws. Transparency, I argued, is a key principle for the protection of privacy; as Elgesem writes, privacy principles ought to be public principles (Elgesem, 1996: 47). Moor too argues that norms of privacy must be public (Moor, 1997: 32). To this, Colin J. Bennett has added the principle of fairness. In 1992, Bennett compared privacy legislation in the US, UK, Germany and Sweden to find that privacy laws were built around six principles of fair information management, including transparency, and also fairness and consent (Bennett, 1992: 95-110). Bennett's first principle is the principle of openness, as discussed above. The second is the "principle of individual access and correction", which enables an individual to check and correct information held about her. The third is the "principle of collection limitation", prescribing that information be collected for one specific, legitimate, justified purpose. The goal here is to ensure relevance. The fourth is the "principle of use limitations", which generally means that information collected for one purpose should not be used for another purpose. The fifth is the "principle of disclosure limitation", which holds that personal data given to one agency shall not be communicated to another without either the individual's consent or legal authority. And the sixth is the "security principle", providing that personal data should be protected by reasonable security safeguards to prevent privacy breaches from other sources. The internet has changed radically since 1992, but Bennett's principles, founded upon the principles of transparency and fairness (and consent), have proven extremely influential. They appear in legal instruments internationally, including in the Australian Privacy Principles, which came into effect in 2014 (OAIC, 2017b).

A second privacy principle, devoted to transparency and fairness, might be worded: **“Individuals, organisations and governments must be fair and transparent in their dealings regarding a person’s privacy, and in particular in their collection and sharing of a person’s private information.”** Here there is significant overlap with the first principle: if companies use your private data without consent, there is a good chance there has been a failure of either fairness or transparency. To link it even more squarely with the conception of consent I have been describing: without fairness and transparency, the possibility of consent has likely not been offered.

Provisions expanding upon and implementing such a principle would catch the Facebook advertising system Beacon, which was launched in 2007 and shared users’ off-Facebook purchase data with other users, but gave no clear way to opt out. After a public outcry, Facebook adjusted the program; clearly, however, Beacon infringed basic principles of transparency and fairness (Christians et al., 2012: 98). Beacon’s breach of the consent principle is even more egregious. Meanwhile, Julie Brill of the US Federal Trade Commission has said data brokers such as Acxiom ought at least to tell people about the data they collect, how they collect it, whom they share it with and how it is used: “We need to figure out what the rules should be as a society” (Singer, 2012). For instance, what are the rules regarding children? Can Acxiom begin building individual profiles at birth? Even earlier? Or should there be a blanket provision, tied into this prescription for fairness, that harvesting children’s data is not allowed? One suggestion is that laws could be passed making it illegal to disclose *anything* a person shares before the age of 18, thus effectively sealing everyone’s digital juvenile records (Schmidt and Cohen, 2013: 67). In 2013, Ireland adopted rules mandating that companies make clear when they are collecting data about individuals’ online activities, by displaying a relevant icon (Dwyer, 2015b: 129). Currently, Acxiom clearly fails to offer the possibility of consent, and so clearly falls foul of both the first and second principles. Meanwhile, there is the growing expectation that developers of mobile phones apps with geolocation features need to be transparent and open about how data is collected, used and shared (Dwyer, 2015b: 132).

Prima facie, the formula of humanity outlaws deception and coercion. As discussed in chapter five, however, this prescription has exceptions. It is right to lie to the killer at the door. It can be right for a policeman to arrest and detain a suspect. Generally, though, deception and coercion are forbidden, and this too should be reflected in consumer-style protections. Indeed, this is precisely the terrain covered by consumer law in various jurisdictions. In the United States, the Federal Trade Commission's Bureau of Consumer Protection targets unfair, deceptive and fraudulent business practices (FTC, 2017). These "fair information practices", as they are known, are employed by the Federal Trade Commission [FTC] and industry self-regulation bodies to set benchmarks of good conduct (Grimmelmann, 2009: 810). They are not binding law; nonetheless, the FTC has imposed penalties including a US\$22.5million fine on Google in 2012 for placing advertising tracking cookies in the browsers of Safari users, despite assurances from Google it would not do so. The fine was imposed not because Google breached fair information practices, but because it breached a previous order in which Google promised to the FTC not to misrepresent its privacy policies to consumers (Tsukayama, 2012). In Europe and Australia, analogous principles *are* enforceable (Grimmelmann, 2009: 810; OJEU, 2016; ALII, 2017a). Under Australia's *Competition and Consumers Act 2010*, broad prohibitions carry financial penalties. One prohibition prescribes that unfair terms are void (ALII, 2017a: ss. 23-28). A contractual term is defined to be unfair if, *inter alia*, "it would cause a significant imbalance in the parties' rights and obligations arising under the contract" (ALII, 2017a: s. 24). Further, all persons are prohibited from engaging in "unconscionable conduct" (ALII, 2017a: ss. 20-22A). Like "unfair", "unconscionable" is a broad, vague notion. It lacks precision and clarity. This, however, is the consumer law's strength, in that it gives judges and arbitrators a wide discretion to penalise contracts and conduct they regard as unethical.

Crucially, in provisions that square with the formula of humanity, Australian law also prohibits "misleading and deceptive conduct". The Act states: "A person must not, in trade or commerce, engage in conduct that is misleading or deceptive or is likely to mislead or deceive" (ALII, 2017a: s. 18). These provisions are enforced by the Australian Competition and Consumer Commission [ACCC], with adverse findings resulting in financial penalties (Ellery, 2016). In 2017, the ACCC signaled that it would be seeking higher penalties, so that fines would

amount to more than just the cost of doing business (Gould and Hirst, 2017). Explicitly and bluntly, the Australian consumer law prohibits unfair terms, and prohibits conduct that is unconscionable, misleading or deceptive. Similar prohibitions could easily and effectively be applied in relation to privacy. Indeed, Australia's privacy commissioner notes that some such consumer-style protections in favour of fairness have been built into the Australian Privacy Principles contained in the *Privacy Act* (Pilgrim, 2015). When it comes to an individual's privacy on the internet, unfair terms in user agreements should be void, just as unconscionable, misleading and deceptive conduct should be illegal, with penalties enforced for breaches. A third privacy principle might provide: **"Individuals, organisations and governments must refrain from conduct that is misleading, deceptive, unfair or unconscionable in their dealings regarding a person's privacy, and in particular in their collection and sharing of a person's private information."** Here again, there is a clear link to the consent principle, which is itself drawn directly from the formula of humanity. No one could possibly consent, it would seem, in the face of misleading, deceptive, unfair or unconscionable dealings. Such a principle would apply in cases of misrepresentation, where data is being collected for one stated purpose, but is also being used for a different, hidden purpose. The principle would apply in the common scenario that occurs when users visit websites with content from a third party, and that third party then tracks the users' browsing across the web (Mayer and Mitchell, 2012: see chapter two).

A fourth privacy principle could serve to defend against coercion. As with deception, where there has been undue pressure upon an individual to give up privacy, that interaction is, *prima facie*, unethical. This fourth principle might be worded: **"Individuals, organisations and governments must refrain from conduct that is coercive in their dealings regarding a person's privacy, and in particular in their collection and sharing of a person's private information."** Again, this links to the consent principle, given that consent contained in the face of coercion is unlikely to be morally justifying. Such coercion could be argued to exist in certain scenarios involving large companies such as Apple, Google or Facebook. Hypothetically, these companies would be able to change their privacy policies in a way that many users do not want, but to which users might feel compelled to agree. If, for instance, many of my work documents are stored with

one company under a privacy agreement that is entirely acceptable to me, but then the company alters its privacy provisions in a dramatic manner which I find unacceptable, then I may feel compelled to accept the new provisions in order to maintain access to my work documents. There may, in this instance, be undue pressure on me to accept new privacy provisions, even if I would rather not.⁶⁴ The case of government surveillance is more clear. Imagine I want to use Google and Facebook. Imagine further that I fully accept the terms and conditions of their privacy policies. My consent to their privacy practices is informed, genuine and morally justifying. It is also satisfactorily iterative. However, I may also be aware that government agencies are collecting all my Google searches and Facebook activities. To this, I do not wish to consent, but what choice do I have? Arguably, there should be redress against government agencies acting coercively.

Admittedly, in both these scenarios the primary issue involves consent. I have not consented to the new provisions in the first case, nor have I consented to government surveillance in the second case. In neither case have I been offered the *possibility* of consent. But coercion must be minimised too. Consumer-style regulation which explicitly targets coercion is one way of identifying and invalidating practices and policies when the pressure to agree becomes unacceptable. If I force you into a compromising pose for a photo which I plan to share on social media, I am using you merely as a means. General legal prescriptions, tied to civil and criminal remedies, can provide the legal muscle to enforce such an ethical position.

Such general prescriptions, founded on a two-tier model of consent drawn from Kant's formula of humanity, require general exceptions. Privacy, as we have seen, cannot be absolute, but exists in the contexts of other rights, obligations and freedoms. The individual's ability to invoke a right of privacy must be balanced against other rights and interests, including the right to free speech. As Gavison writes: "It is obvious that privacy will have to give way, at times, to important interests in law enforcement, freedom of expression, research and verification of data. The result is limits on the scope of legal protection of privacy" (Gavison, 1980: 457). In general terms, the right to privacy must also be balanced against the public interest (e.g. ALRC, 2014: 7-8). Indeed, a public interest test is

⁶⁴ Arguably, this is a case of exploitation rather than coercion. It is a fine line between the two. I use this example because it is a hard one; more obvious cases of coercion are easier.

particularly compelling. In specific cases, for instance, a government might engage in coercion to obtain private information. If there is a well-founded suspicion that a serious crime is about to be committed, a person might justifiably be compelled to hand over private computer files. (For a related discussion, see the conclusion to chapter five, above.) A balancing act must be struck between, on the one hand, people's right to privacy and, on the other hand, the rule of law and people's right to security, *inter alia*. Invoking the public interest can help to strike that balance. General principles that clearly articulate the values and rights at stake and the need for such a balancing act with the right to privacy could then be interpreted and applied by an independent judiciary. A fifth privacy principle, then, would provide: **"The right to privacy is not absolute, but exists in the contexts of other rights, obligations and freedoms. These include, but are not limited to: the right to personal security; the right to be connected; the right to free speech; freedom of the press; and freedom of information. The right to privacy needs to be balanced against these other rights, obligations and freedoms. This balancing act must take into account the public interest."** As a liberal communitarian, Etzioni argues that individual rights have the same fundamental standing as the common good, and it is up to each society to determine between these two claims. One society may value individual rights more highly; another may value the common good (or public interest) more highly. Hence a balance must be struck (Etzioni, 2015: 1271-1272). So it is with privacy. Consider a politician having an affair. Would it be an invasion of privacy to reveal her liaison? In the US, the law of "public figures" ensures their right to privacy is much more limited than that of the ordinary person (Whitman, 2004: 1196). In the UK and Europe, by contrast, private matters of public figures are generally considered off limits, by the law as well as by etiquette (Bennett, 1992: 83, n. 96; Whitman, 2004: 1170). As Etzioni argues, different societies will approach the balancing act differently. Various positions are defensible, though I propose, in line with principle five, that such revelations are only justified if they serve the public interest (such as if the affair is likely to compromise the politician in the exercise of her public duties).⁶⁵

⁶⁵ In the US, a test is indeed applied to determine if there is a legitimate public interest in a private revelation. However, courts have generally defined the public interest broadly, arguably too broadly, thus allowing publication of highly personal details (Whitman, 2004: 1196).

Let us return, then, to the surveillance programs of government agencies such as the NSA. Seeking to counteract terrorism by surveilling its citizens, the US government was seemingly acting on the principle that the individual's right to privacy was outweighed by the individual's right to personal security. Above, I noted that a law only conforms to Kant's united will of the people if it is a *just* law, and it is a just law only if it is a law to which the people could have assented. For the sake of argument, let us assume again that the NSA was operating entirely legally under wholly just laws. (A big assumption, granted.) *Prima facie*, the NSA has breached many of our privacy principles. It did not obtain the people's consent. It was not behaving with transparency. Its clandestine operations, involving spyware and partnerships with corporations, were misleading and deceptive. The crucial privacy principle thus becomes the fifth, which demands the consideration of other rights and the public interest. Here, did the right of citizens to safety outweigh the right of citizens to privacy? This is a matter open to dispute. Some researchers argue that targeted surveillance is just as effective as blanket surveillance in countering terrorism: "Mass surveillance makes the job of the security services more difficult and the rest of us less secure" (Corrigan, 2015). Blanket surveillance conducted in secret is, it seems, a tremendous invasion of the right to privacy with questionable benefit to personal safety. What's more, as I argued earlier, on the question of just laws there seems to be no compelling argument for the NSA to devise and implement the broad sweep of its programs *in secret*. Programs such as PRISM and KEYSTROKE sought to serve the public interest, certainly, but the public was never given a chance to discuss whether or not an appropriate balance had been struck between this public interest and the right of privacy. As Lessig writes, policy makers must always ask what mix of law and technology will restore the proper level of user control: "That level must balance private and public interests" (Lessig, 2006: 200).⁶⁶ Conducted in secret, NSA surveillance (and analogous surveillance in other countries) was implemented to protect people's freedom, but only did so by limiting people's freedom. For a proper balance to be struck, public debate was needed in the first

⁶⁶ I support Lessig's analysis, though I baulk at the polarisation of private interests and public interests. As I have been arguing, privacy is also a public interest. It plays an indispensable role in connecting us to others, and only operates for people as embedded social beings, rather than as idealised individuals. This is one reason why the public interest does not warrant its own principle as a limiter of privacy, but is rather one ingredient in the fifth principle, which also accounts for rights and interests including free speech, freedom of information, freedom of the press, and so on. Nonetheless, I endorse Lessig's overarching point that both law and code must play a part in effectively balancing competing interests. I explore this point below, in section three.

place. Then, once such laws and policies are enacted, remedies based on consumer law can help. Harnessing both civil and criminal remedies, privacy protections based on consumer law have tremendous potential to parry improper challenges from governments, as they do for challenges from individuals and companies.

Note, however, that I am not advocating that privacy protections ought to be enacted to *duplicate* consumer protections. As Grimmelmann argues:

I am not calling for the direct application of products liability law to online privacy ... Instead, I am suggesting a process of thoughtful conversation and translation between two bodies of law that have a common history and more in common than scholars and lawyers sometimes realize (Grimmelmann, 2009: 826-827).

It is the *approach* that I am advocating. This approach involves general prescriptions in line with Kant's formula of humanity. Kant's formula gives us our consent principle, which involves a model of actual consent defined also to include possible consent. This application of individual consent must then be supplemented by collective consent in the shape of just laws, in the same way as Kant's ethics are supplemented by Kant's political and social philosophy. These laws must abide by morality and, where appropriate, reinforce individual consent. Hence I propose laws that mirror consumer protections, drafted so as to embody the principles given above, all of which are in line with the formula of humanity. Of course, collective consent will no doubt be flawed in practice, given the limits of legislatures. None of us is perfectly rational individually, much less so collectively. Nonetheless, the model I have sketched out gives us a firm normative grounding. In line with the five principles articulated above, we can formulate laws modeled on consumer protections that would help significantly to prevent and remedy privacy abuses from government agencies, social media companies and more. With enforceable provisions and sufficient penalties, such laws would be potent. If researchers want access to anonymised medical data, the law could spell out what (if any) individual consent is required, and under which conditions such data may be used. If a cloud-based data storage company wants to alter its terms and conditions regarding privacy, the law would insist that such alterations are fair, transparent and in the spirit of the provisions originally consented to. And if a government agency wants to engage in surveillance, the law can set limits, including by mandating transparency and fairness, except in clearly delineated

exceptional circumstances. Further, these laws ought to involve a multi-pronged approach. They might involve a regulatory body, and perhaps an ombudsman or commissioner, to receive complaints and resolve disputes; they ought to involve civil remedies, including (but definitely not limited to) the enactment of a tort of serious invasion of privacy, as discussed in section one; and they might involve criminal penalties.

A final point. I have argued that privacy protections ought take their cue from consumer law protections. This is not to suggest, however, that internet users are to be *treated as consumers*. Rather, internet users are to be regarded as *persons*, and privacy protections could be drafted in a manner analogous to consumer law protections. The distinction is significant. Throughout this thesis, I have been talking of persons, individuals and citizens. The concept of an internet user that I have been trying to sketch is in a Kantian sense: of a rational being with dignity and autonomy. Further, in chapter five, I suggested that accounts of individual privacy must recognise that individuals exist not in isolation, but as beings-in-relation. Privacy, I noted, only has meaning in relation to friends, family, community, society, and more. Privacy is one way in which individual and group distinguish themselves from another, but also one way in which people create the ties that bind, enabling individuals to coalesce into groups. Privacy can connect just as it can separate, leading me to propose the notion of relational privacy, in which people are not merely consumers, but agents with relational autonomy. Hence I am talking not about the privacy of consumers, but of persons. This, in turn, has a profound effect on the remedies and protections required for privacy. As Sarigol *et al.* write in their study of shadow profiles: “... we should consider privacy as a collective concept, where individual privacy policies are not sufficient to control private information” (Sarigol et al., 2014: 104). Remedies based on consumer law, I suggest, would go some way towards acknowledging and protecting a privacy that is relational.

III - Beyond consent: extra-legal protections

Sometimes, however, law is irrelevant. Privacy infractions come in many hues, and many are fleeting, unnoticed or inconsequential. As Grimmelmann writes:

Many privacy harms, embarrassing though they may be, are beneath the threshold at which the law ought to take notice. The fact that your mother found out your plans to attend International Skip School Day is not, and should not be, a legally cognizable harm (Grimmelmann, 2009: 808).

Similarly, Gavison writes: “The law, as one of the most public mechanisms society has developed, is completely out of place in most of the contexts in which privacy is deemed valuable” (Gavison, 1980: 459). For many privacy breaches, the collective consent of the law has no role to play.

What’s more, sometimes individual consent has no role to play either. As I showed in chapter five, the formula of humanity is not satisfied by an adherence to consent alone. The principle of consent derives from the mere means principle, but the formula also contains the ends in themselves principle, which commands that we always treat others, and ourselves, as ends in themselves. One way in which the ends in themselves principle often manifests is in our personal relationships. In chapter four, I argued that one compelling justification for privacy is that without privacy we cannot trust, love and befriend. This argument seems particularly apt for the internet, where social media, dating platforms, hookup apps and adultery sites are radically rewriting the way we meet and mate. One result is that the very notion of friendship is becoming more varied and dynamic (Goggin and Crawford, 2010). This may not be a good thing. Sherry Turkle argues that our internet interactions are jeopardising our friendships, and indeed our ability to befriend, partly by their encroachments on privacy (Turkle, 2011: 293, 345). It is beyond the scope of this thesis to investigate this claim. Rather, in this section I will be looking at extra-legal means by which privacy can be supported and protected. Sometimes these extra-legal means *do* concern consent: they sometimes re-empower individual consent; and they sometimes complement and support the collective consent of the law. However, sometimes these extra-legal means are suited to operate in cases when individual consent is irrelevant, or when the collective consent of the law ought to remain silent. In some cases, these extra-legal means can provide ways to foster the ends in themselves principle, a principle that, ultimately, is about users treating one another (including our friends) with respect.

The three extra-legal means I will address are social norms, market forces and coding. In this, I follow Lawrence Lessig, who identifies four regulatory “modalities” that affect people’s behaviour: law; norms; market; and architecture (Lessig, 2006: 121-125). They work together, and sometimes in contradiction, to influence behaviour. “The constraints are distinct, yet they are plainly interdependent. Each can support or oppose the others” (Lessig, 2006: 124). The four modalities are particularly evident on the internet. What follows, then, is recognition that, even if our goal is to apply the formula of humanity to internet privacy, a two-tier model of consent is not enough. In order to encourage behaviour that protects privacy appropriately, which includes taking into account privacy’s effect on relationships, we need to look beyond consent and the law. In some cases, this involves calling upon the ends in themselves principle of the formula of humanity, which exhorts us to follow the positive, imperfect duty of treating all others (and ourselves) always as persons, as autonomous agents pursuing their own projects.

i. Social norms

Previously, I have been arguing that significant legal regulation is required to protect privacy in order to reinforce, limit or otherwise affect collective consent. This follows from an application of the formula of humanity as, for Kant, the united will of the people can take shape in the form of just laws. Indeed, the united will of the people can *only* take shape in the form of just laws. For Kant, social norms and other behaviour-affecting factors can never constitute a binding expression of the united will of the people. Social norms can never pronounce, definitively, what is required, permitted or forbidden. However, given that the legislature has recourse to social norms when it enacts legislation, we might think of social norms as potential statements of the united will, and hence *potential* just laws. What’s more, we can certainly acknowledge that norms have a dramatic effect on behaviour. In recent decades, notes Etzioni, scholars have rediscovered the significance of social norms (Etzioni, 2000: 157). Unfortunately, norms are often unclear online (see chapter two). How am I supposed to behave in the “walled garden” of Facebook? On the anonymous bulletin board 4chan? In the swipe right, swipe left dating pool of Tinder? Elsewhere, interactions have been stripped of familiar cues and contexts that help confer guidance and meaning. On

the musical.ly app, the line between lip syncing and soft porn becomes blurred when teenagers engage in highly sexualised dance moves (Munro, 2016). In a Sydney courtroom, one witness nearly caused a murder trial to be aborted by sending a Facebook friend request to a juror (Rigney, 2017). The mediation of the internet can foster distance and depersonalisation, creating the illusion of a norm-free environment. However, our internet interactions are not just virtual, but *real*, and norms do apply.

Which social norms *ought* to apply to internet privacy? One proposal involves harnessing the internet's interactivity by blending crowdsourcing and machine learning to enable the automated "discovery" of privacy norms. Following the contextual integrity framework, researchers can "elicit informational norms based on a crowdsourcing approach that queries users on their privacy expectations based on automatically generated privacy statements" (Shvartzshnaider et al., 2016: 1-2). To identify or prescribe social norms is beyond the scope of this thesis; however, any prescription ought rely heavily on the notion of respect, as commanded by the ends in themselves principle. "Respect for humanity ... underlies all social relationships as a normative ground" (Anderson-Gold, 2010: 28). Ideally, these agreed-upon norms would yield a set of guidelines. Already, Brazil has codified such a framework. In 2014, Brazil enacted an internet code of conduct, the Marco Civil da Internet, with Article 1 providing: "This Law establishes principles, guarantees, rights and obligations for the use of the internet in Brazil ... " Article 3 then names the "protection of privacy" as a fundamental principle (Medeiros and Bygrave, 2015). To a significant degree, a set of guidelines for internet behaviour might overlap with the consumer protections outlined above. Also, they should be expressed in lay language, not legalese. What's more, these guidelines should be drafted in at least two forms: succinct and extended. As with the notice-and-consent model I have been advocating, this would enable their use in an appropriate and iterative manner. The succinct version would suit situations in which users have limited time, attention or screen size, such as when using a smartphone; the latter would suit situations where users have more time, attention and screen-size, such as on a desktop.

Assuming we can agree upon a desirable set of norms and articulate them in a set of guidelines, how can positive social norms be encouraged? The answer, I

suggest, lies largely in habituation and education, which are crucial for the development of virtue. For Kant, virtue does not lie in actions performed merely by habit, as if by rote; rather, it lies in habits and aptitudes that have been acquired by “considered, firm, and continually purified principles”. Only in this way will our virtue be sufficiently deft to deal with novel situations and new temptations (Kant, 1996a: 6:383-384; Herman, 1993: 78; Sherman, 1997: 161). This is one of the reasons we have a duty of friendship, wrote Kant: friends can reflect ourselves back to us, making us better able to see our moral intention more clearly (Kant, 1996a: 6:392). In her Kantian defence of privacy, Sharon Anderson-Gold argues that we can promote respectful behaviours in others through our own respectful behaviour. One way to do this is to contribute to the moral development of others, including future generations, through the support of good manners in our social interactions (Anderson-Gold, 2010: 28). These arguments apply neatly to the net: if the norms we bring are respectful and polite, we encourage this in others.

Habituation, education and setting a good example are especially important when it comes to children. Compared to adults, children are in a state of moral (and neural) plasticity. Online, where norms can be unclear, they are at particular risk of compromising their own and others’ privacy. On Barbara Herman’s Kantian account, the moral *sensitivity* of children must be cultivated. As part of their socialisation, children must learn “rules of moral salience”, akin to an early warning system for moral danger, which can then help provide a practical framework within which to act.

When the rules of moral salience are well internalized, they cause the agent to be aware of and attentive to the significance of ‘moral danger’ ... The rules of moral salience constitute the structure of moral sensitivity (Herman, 1993: 78; see section four, below).

Education can make primary and secondary school-aged children sensitive to moral danger by explaining appropriate privacy norms, both for themselves, and for others. In chapter two, I outlined research showing the extent to which young people share explicit images and videos (Lenhart, 2009; Strassberg et al., 2013). Particularly for sexting, desirable norms ought to be articulated in education and media literacy programs. Researchers suggest that such norms ought to avoid victim blaming and abstinence-only advice, recognising that sexting involves a careful negotiation between risks (including to privacy) and benefits for identity, intimacy, sociability and, ultimately, trust (Hasinoff and Shepherd, 2014: 2949-

2950). Researchers have further recommended age appropriate education in schools about media material concerning love, sexuality, gender and relationships, and that parents be trained to talk to children about the information and values they take from the media they consume (Lumby and Albury, 2010: 151). In Australia, children can study to obtain their eSmart Digital Licence, which is intended to be incorporated into the school curriculum “to prepare Australian children (aged ten and over) to be smart, safe and responsible digital citizens” (eSmart, 2017). In the US, similar initiatives are in place to promote the digital literacy of children, as well as teachers (CommonSenseMedia, 2009; DigitalLiteracy, 2017; Edutopia, 2015).

A vast body of scholarship exists exploring diverse approaches to digital literacy education (eg. Eshet-Alkalai, 2004; Hobbs and Jensen, 2013). There is, however, limited scholarly evidence linking an increase in media literacy with a reduction in harm (Livingstone and Hargrave, 2006: 42). What’s more, it has been argued that education and “empowerment” are less effective than regulation (Rush, 2012: 167-168). Indeed, as I have been arguing, regulation is crucial. Hence I propose digital literacy education as a supplement to regulation, for children as for adults (Meese, 2015: 144-145). For instance, with many online interactions, users are unclear about how their data is used. Education can help to validate consent: if a user is properly informed about potential data uses, then it is more likely there has been a genuine offer of possible consent. For instance, it has been argued that the only way for users to be autonomous in the information society, where internet profiling is often conducted without consent, is via digital privacy literacy, which will enable users to understand profiling and thus to interact with their own profiles (Degeling and Herrmann, 2016). Consent must be a key element of any digital literacy program. What constitutes morally justifying consent? When is consent required? When is re-consent required? In the face of a prevailing sentiment in some quarters that “privacy is dead”, such education ought not just to reveal the *significance* of consent, but also its *significant limits*. Hence such education ought also reveal the role of the law in establishing the right of privacy. Further, such education is needed not just for users, but for companies and governments, so that they too are aware of the importance and limits of individual and collective consent. In this way, pessimistic resignation might evolve into determined involvement.

In 2012, charity worker Lindsey Stone posed for a photo at the Tomb of the Unknown Soldier in the United States. In front of a sign demanding “SILENCE AND RESPECT”, Stone mimed yelling and raised her middle finger. In an analogue world, the photo might have bemused close friends; on Facebook, where Stone’s posts were set to public by default (a point addressed below), the image spawned a “Fire Lindsey Stone” Facebook page which attracted 19,000 likes. Stone was indeed fired, becoming one of many whose lives have been upended by public shamings (Ronson, 2015: 253-254). However, the potential also exists for public laudings. There are, I have been suggesting, a series of ways in which to shape social norms that regulate behaviour concerning internet privacy. These include the articulation of privacy-protecting guidelines, digital literacy education (and habituation) of children and adults, and even the practice of positive feedback for privacy-respecting online behaviour. No doubt this list could be complemented by other norm-shifting initiatives. Lessig, for instance, argues that norms among commercial entities could be shifted to help build *trust* around specific privacy protective practices (Lessig, 2006: 223). Social norms are significant drivers of behaviour. If social norms come to embody a greater value for privacy, then spyware, Creepshots and the on-sharing of sexts will be more widely and reflexively recognised as the serious, harmful invasions they are.

ii. The market

Another of Lessig’s four regulatory modalities is the market. Indeed, a company such as Facebook can be subject to competing imperatives: on the one hand, the business imperative to obtain as much data about as many users as possible; on the other hand, the moral imperative to respect the privacy wishes of those users, which may be strict. Business imperatives can thus compete with moral imperatives (Spinello, 2011: 42-43; Christians et al., 2012: 97-99).

One proposal to align business imperatives with moral imperatives is to link the compensation of executives in social network systems directly to the privacy protections they provide (Helman and Hannes, 2016). We currently have a behavioural market failure, the authors argue, where many social media users act against their own best interests, and where social media firms have no incentive to

internalise the privacy interests of users. Their proposal is to factor data management practices into executive pay by giving companies an annual privacy rating based on technological measures and user satisfaction. Pay rates would then be assessed with the help of a compensation committee. This, they contend, would inject privacy competition into the market (Helman and Hannes, 2016). As I have shown, privacy is valuable for many reasons, including the non-instrumental justification of dignity. As such, there are potential risks with putting a dollar value on its protection (Sandel, 2012). However, as one of four regulatory modalities, market forces could justifiably and effectively prompt companies to respect privacy. Indeed, Schmidt and Cohen argue that technology companies will increasingly find themselves beset by public concerns over privacy, and would thus be wise to take proactive steps, including: offering a digital “eject button” that liberates all of a user’s data from a given platform; not selling personally identifying information to third parties or advertisers; or perhaps even not selling *any* data to third parties. A group of companies might band together and make a pledge to abide by such steps (Schmidt and Cohen, 2013: 66-67). Like social norms, market forces can potentially be shaped and shifted to protect privacy more effectively.

iii. Coding: privacy by design

Finally, Lessig highlights the crucial regulatory power of architecture. Those who want to regulate the web, he writes, must pay attention to the precise way in which coders are creating apps, platforms, websites, operating systems and other types of software and hardware. Proposing that code is law, Lessig writes that we can “build, or architect, or *code* cyberspace” either to protect or to doom the values we take to be fundamental. Code isn’t found, but made, “and only ever made by us” (Lessig, 2006: 5-6).

As discussed in chapter two, the values embedded in the internet’s coding have changed dramatically since 1969:

The invisible hand of cyberspace is building an architecture that is quite the opposite of its architecture at its birth. This invisible hand, pushed by government and by commerce, is constructing an architecture that will perfect control and make highly efficient regulation possible. The struggle in that world will not be the government’s. It will be to assure that essential liberties are preserved in this environment of perfect control (Lessig, 2006: 4)

Internet coding, Lessig writes, has tended to underprotect privacy. Indeed, it continues to do so, as privacy comes under threat from government surveillance and commercial data-gathering. Hence Lessig recommends “privacy enhancing technologies”: “These are technologies designed to give the user more technical control over data associated with him or her” (Lessig, 2006: 223; see also Bambauer, 2013). The notion of “technical control” clearly aligns with consent; hence coding can be a way to enforce the mere means principle. Coding solutions include putting locks on photos and also ephemeral messages, which disappear after a short time (such as those sent via Snapchat) (Hasinoff and Shepherd, 2014: 2945-2947).

One general privacy-enhancing technology is encryption, which both protects privacy and enhances security (see chapter two). One version is strong-encryption-by-default, which can protect against hackers aiming to steal personal information, companies eager to profit from personal data and (perhaps) governments bent on surveillance (Levy, 2016b). In 2016, Facebook introduced an opt-in system of encryption for its Messenger instant messaging service (Yadron, 2016). That same year, whistleblower Edward Snowden derided Google’s new “Allo”, a chat app with a virtual assistant, which did not feature end-to-end encryption as, “A Google app that records every message you ever send and makes it available to police upon request” (Hackett, 2016). There are loud arguments for and against: privacy advocates argue it’s needed for individuals to communicate freely; government agencies argue that any encrypted systems should be designed with a back door, in the fight against terrorism and other crime. The issue surfaces sporadically, including in 2016, when Apple resisted calls from the US government to develop code to unencrypt its iPhone (Levy, 2016b). Debates surrounding encryption are likely to become more heated with the advent of quantum entanglement, which promises eavesdrop-proof communication (Wen, 2016: see chapter two). Certainly, encryption is no panacea. As Schneier writes: “My guess is that most encryption products from large US companies have NSA-friendly back doors” (Schneier, 2013).

In the debate about encryption lies the realisation that privacy must be balanced against other rights and freedoms, as addressed in principle five of our consumer-

style protections. Indeed, internet companies recognise that coding (and technology generally) plays an important role in striking a balance between various interests and rights. On the one hand, Google “has always operated on the belief that more access to information generally means more choice, more power, more economic opportunity and more freedom for people” (Yorke, 2015). On the other hand, the company says it values users’ privacy: “Respecting user privacy is both a moral and business imperative for Google” (Yorke, 2015). To this end, all Google engineers complete mandatory privacy training, all Google products are regularly reviewed by privacy experts, and Secure Sockets Layer [SSL] encryption, which encrypts data between server and browser, operates as standard to prevent others “snooping” on a user’s activity when on an open network, such as using a laptop at a café (Yorke, 2015: see chapter two). At times, Google has paid the price for developing products that inadequately protect user privacy. This, as discussed in chapter two, was arguably behind the failure of Google Glass, which enabled (with a simple hack) users secretly to photograph and film others (Edwards, 2016). By contrast, a subsequent pair of internet-enabled glasses, called “Specs” and developed by the company behind social media platform Snapchat, alert others with an illuminated light whenever the user is taking photos or filming. As one commentator wrote:

It’s privacy by design. The coolest part of the Specs isn’t what they look like, it’s how they show respect for privacy. The moment you begin recording the world knows about it, via a light mounted in the frame of the specs that lets people know they’re being filmed (Edwards, 2016).

Specs store captured video on the device. Google Glass, by contrast, stored the footage in its cloud servers, where Google owns the data and can use it to build user profiles (Edwards, 2016). As Lessig notes, privacy protections *can* be engineered into software and hardware. Such engineering, argues Waters, can empower individuals to give informed consent: “The technology, the architecture of the internet and apps offers an enormous potential for increasing the ability of individuals to make informed decisions” (Waters, 2015).

In chapter two, I noted that privacy has only been coded into the net as an afterthought, via add-ons and browser extensions. Slowly, however, such moves are gathering pace, captured in the phrase “privacy by design”. In Europe, privacy by design is a fundamental principle of the GDPR. Article 25, “Data protection by

design and by default”, provides for the implementation of “appropriate technical and organisational measures, such as pseudonymisation, which are designed to implement data-protection principles ...” (OJEU, 2016: 48). This enables two regulatory modalities, the law and coding, to work together to provide a strong protection of privacy. In this way, the GDPR will mandate the type of architecture that researchers have recommended in more specific cases. In 2011, it was argued that there should be no “publicly” available fields on Facebook unless the user explicitly chooses otherwise (Spinello, 2011: 45). The GDPR will go far towards changing the industry standard to ensure that information is private by default, and public by choice. This alone will potentially make a dramatic difference. Such regulation may have saved Lindsey Stone’s job and reputation.

More inventive possibilities exist. Funded by the Finnish government, one initiative aims to strengthen the role of individual consent by establishing an information repository that serves as a go-between. On this repository, known as a “MyData” account, a user can see and decide what information is held about her and who has access to it. Personal data is thus controlled from one place even as it is created, stored and processed by hundreds of different services. Hence the flow of consents is distinguished from the flow of data.

The aim is to provide individuals with the practical means to access, obtain, and use datasets containing their personal information, such as purchasing data, traffic data, telecommunications data, medical records, financial information and data derived from various online services and to encourage organizations holding personal data to give individuals control over this data, extending beyond their minimum legal requirements to do so (Poikola et al., 2010: 3).

The benefits, arguably, would be many. Among them: individuals would have transparent data management tools; opportunities would arise for new data-based businesses; and companies and governments would have certainty, clarity and good outcomes. Internet privacy, according to the researchers, would become “human-centric” (Poikola et al., 2010: 2).

Coding can empower consent and dissent in other ways too. In chapter two, I discussed the guerrilla tactics being employed by some internet users to protect their privacy. These tactics can include the “obfuscation” offered by the tracker-blocker Disconnect Me and the TrackMeNot browser extension, which confuses trackers by hiding actual web searches within fake web searches (Nissenbaum and

Brunton, 2015; Disconnect, 2016; TrackMeNot, 2016). They include hampering companies' attempts to create user profiles by employing an "informed dummy generation strategy" (Degeling and Herrmann, 2016). Privacy-protecting alternatives are proliferating: Google has a rival in search engine DuckDuckGo; and Facebook's data harvesting inspired the social network Diaspora (Nissenbaum and Brunton, 2015; Christians et al., 2012: 99; see chapter two). Using TOR makes you considerably more anonymous (see chapters one and two). As we have seen, software and hardware have embedded values. Via coding, significant scope exists either to protect privacy, or to undermine privacy.

In cases of trust, love and friendship, it may well be the extra-legal measures of social norms, market forces and coding that have the best chance of protecting privacy. These extra-legal measures might encourage me not to post a photo of my wife to social media that might harm her friendships, cause her to be fired, and so on. Erosions of privacy threaten our friendships and our ability to trust and to love. Apart from forbidding us from using one another merely as a means, Kant's formula of humanity also mandates respect and love. Social norms, market forces and coding can reinforce the limits set by individual and collective consent, certainly; but social norms, market forces and coding can also do more than that if, by encouraging respect, they prompt internet users always to treat everyone, including themselves, as people who set their own ends. Above all, if all four regulatory modalities work together, privacy has a good chance of being well protected.

IV – Welcome to *cosmoikopolis*

Even in the face of a four-pronged regulatory approach, the internet presents a challenge, because it transcends the national. In some ways, it transcends the international, as a network that exists in the virtual beyond. This presents a significant challenge for all four regulatory modalities, and in particular for any attempt to impose *legal* regulation. Hackers tend not to be bothered by border guards. Google and Facebook are globe-straddling multinationals. The NSA in the United States not only collects data worldwide, but exchanges data with the ASD

in Australia and GCHQ in England, not to mention a long list of other government agencies at home and abroad (Schneier, 2015: 63-72; see chapter two).

Drawing on Kant's formula of humanity, I have built my core argument upon consent, both individual and collective. But how can consent possibly be effective when the medium is so comprehensively international? Let us again assume that the NSA has the collective consent it requires to be morally justified in its surveillance within the USA. However, an internet user in Australia is also subject to NSA surveillance, as are users in Israel, Ireland and Iceland. Everyone, everywhere, it seems, is significantly affected by the NSA's activities, at least potentially. Does this mean a law enabling NSA-style surveillance can only be just if it has the possible consent of *everyone in the world*? It seems so. What's more, the NSA reportedly treats non-US citizens as inferior to US citizens. If this is so, then the laws under which it is doing so are necessarily unjust, just like Kant's example of a law authorising hereditary privilege. Citizens globally could not possibly consent to such a law. However, assuming that the NSA is treating all citizens of the world equally, and none merely as a means, then international laws could settle the matter definitively. If there are *just* international laws (that treat all the world's citizens equally), and the NSA is operating within such laws, then its activities are justified. If there is no international law, the NSA must instead work within the parameters set by the law of each land. In any country in which it operates, it must obey the law, which must be just. This then means that the NSA, and also Facebook and individuals, must obey a variety of standards from country to country, according to legal provisions. This makes for an onerous hotchpotch of obligations. A better approach involves a coordination of international and national protocols, and at times transnational and local protocols.

Such an approach, says privacy advocate Nigel Waters, makes the regulation of internet privacy difficult, but not impossible:

It has been and will continue to be a constant struggle for both domestic and international law to deal particularly with online privacy, given the nature of the media that it is, but I don't think that's an argument for giving up or being fatalistic about the ability to have an appropriate framework. I think there have been some surprisingly well developed attempts at cross-border regulation, providing frameworks that can deal with the essential open nature of the internet (Waters, 2015).

As we have seen, international instruments already play a pivotal role. Currently, privacy is protected by the International Covenant on Civil and Political Rights (ICCPR), which came into effect in 1976 and provides, in Article 17, that: “No one shall be subjected to arbitrary or unlawful interference with his privacy, family, home or correspondence, nor to unlawful attacks on his honour and reputation” (OHCHR, 1976). As at February 2017, the 169 countries that have ratified the ICCPR include the US, UK, Australia and European countries such as Germany and France. Six countries have signed but not ratified the covenant (including China) and 22 countries have not signed (including Saudi Arabia, Malaysia and Myanmar) (OHCHR, 2017). The ICCPR follows on from Article 12 of the United Nations’ Universal Declaration of Human Rights (UN, 1948). At a transnational level, the Paris-based Organization for Economic Cooperation and Development adopted eight privacy principles in its 1980 OECD Guidelines. The first provides: “There should be limits to the collection of personal data and any such data should be obtained by lawful and fair means and, where appropriate, with the knowledge or consent of the data subject” (OECD, 1980). From May 2018, the European Union will support the OECD privacy principles in its GDPR, which enacts strict standards for consent, the transparent processing of data and the right to erasure (OJEU, 2016: see above). In Article 51, the GDPR stipulates that each member state shall establish its own independent public authority to “contribute to the consistent application of this Regulation throughout the Union”; and, under Article 60, each member state’s supervisory authority is directed to cooperate with a lead supervisory authority.

There is thus one level of privacy regulation that is international and transnational, and that transcends individual nation states, just as the internet does. Meanwhile, each country must enact its own particular privacy provisions. In practice, jurisdictions already operate with a blend of international and national regulation. Australia, for instance, has ratified the ICCPR and has passed laws including the federal *Privacy Act 1988* (see chapter three and section one, above). The Australian government has also sought to cooperate with other governments. As a spokesperson for the federal Attorney-General’s Department says:

Australia is a member of the Asia-Pacific Economic Cooperation forum’s Data Privacy Sub Group, which developed a cross-border privacy framework that is the blueprint for greater regional cooperation on privacy rules and enforcement (Attorney-General, 2016).

Unfortunately, no single country seems to have attained a perfect blend of international and national regulation, even if some jurisdictions are nearer the mark than others: “I think possibly some of the Canadian jurisdictions and New Zealand are probably as close as you get to a model but even they have got some significant flaws” (Waters, 2015).

International law alone is unlikely to suffice. National laws have an important benefit: they allow for variation from country to country. Hence the law of privacy can be tailored to the history and values of a particular populace. As discussed in chapter four, every known society values and respects privacy, but privacy norms vary dramatically over time and between cultures. Hull *et al.* argue: “There is no such thing as a universal privacy norm” (Hull et al., 2011: 290). My theoretical starting point has been Kant’s formula of humanity, which applies universally. The dual principles of always treating others as ends in themselves, and never simply as means, allow for no exception. Nonetheless, we can allow for regional variation, as long as variations conform to these stringent dual prescriptions. The principle of never using anyone merely as a means can be observed by the implementation of a two-tier model of consent, supplemented by extra-legal measures; the principle of always treating people as ends in themselves involves a more open-ended application of respect, which often requires extra-legal measures. On a Kantian model, the categorical imperative must apply universally and without exception, but that does not mean it cannot adapt to local circumstances, customs and social norms. Above, I noted Barbara Herman’s rules of moral salience, or RMS, which can help identify situations when moral judgment is needed. These rules, she writes, are not fixed, but represent the prevailing moral climate:

What is attractive about introducing RMS into a Kantian theory of moral judgment is that it would seem to let us have it both ways: while morality has an objective foundation, we have good *positive* reason to tolerate *some* culturally based moral differences” (Herman, 1993: 92-93).

Sexting can be tolerated, but misogynistic spyware cannot.

With the formula of humanity, Kant prescribed a normative ideal for individuals. Kant also prescribed a normative ideal for a world in which the categorical imperative is universally observed. For Kant, this ideal state of perfect morality is

a “realm of ends”, where people only ever treat one another, and themselves, as ends in themselves, and never merely as means (Kant, 2009: 433-434). A realm of ends can prevail only if privacy is respected in such a way that the formula of humanity is universally observed. While the realm of ends is merely an *ideal*, Kant also prescribed a multi-faceted moral and political *goal* in the shape of cosmopolitanism, or world citizenship. The notion has its origins in Classical Greece, when Diogenes declared his allegiance not to the *polis*, but to the *cosmopolis* (Kleingeld and Brown, 2014). For Kant, who argued that all individuals are self-legislating members of a universal moral community, cosmopolitanism was a perfect fit (Kleingeld and Brown, 2014). At the level of the state, Kant wrote, we find a mirror of individual morality: just as people are to be governed by the internal legislation of morality and the external legislation of just laws, so too states must be internally organised and externally organised in a manner consistent with peace (Kant, 1996c; Kleingeld, 2014: 66-68). Further, wrote Kant, all states must recognise the rights not only of their own citizens, but of the citizens of other countries (Kant, 1996c). Only then can states come together in a way that will bring humanity “ever closer to a cosmopolitan constitution” (Kant, 1996c: 8:358). “Only in a universal *association of states* (analogous to that by which a people becomes a state) can rights come to hold *conclusively* and a true *condition of peace* come about” (Kant, 1996a: 6:350). Kant is, finally, optimistic about human nature and ethical progress:

In this way, remote parts of the world can establish relations peacefully with one another, relations which ultimately become regulated by public laws and can thus finally bring the human species ever closer to a cosmopolitan constitution (Kant, 1996c: 8:358).

In this future world envisioned by Kant, cosmopolitan right prevails.

More than 200 years after Kant’s death, cosmopolitanism has indeed been realised, at least to a degree. The United Nations, founded in the wake of World War II, is a septuagenarian realisation of the notion that people belong to a single world order bound by moral universals. Presumably, its existence would have thrilled Kant (although presumably he would have been critical of its failures too). Instruments such as the Universal Declaration of Human Rights and the ICCPR, complete with privacy protections, comprise cosmopolitan law. Meanwhile, a quasi-cosmopolitan approach is embodied in transnational protocols such as the GDPR. To some extent, a cosmopolitan right of privacy has been articulated and

protected. And the internet, I propose, is perfectly suited to such a cosmopolitan approach. As it tends increasingly to convergence, ubiquity and multi-directionality, the internet is remarkable for its ability to cross borders and boundaries. Recently, the internet's global nature has led to calls for an international bill of rights. In 2014, Tim Berners-Lee called for an online Magna Carta to protect the rights of users worldwide, partly in the face of increasing threats from corporations and governments (Kiss, 2014; see Schneier, 2015: 210-212). And in July 2016 the United Nations took a step towards enacting an international Magna Carta for the digital age by adopting a resolution stressing "the importance of applying a comprehensive human rights-based approach when providing and expanding access to the internet and for the internet to be open, accessible and nurtured by multi-stakeholder participation." The UN declared, in short, that "online freedom" is a human right (UNGA, 2016). I will not argue here for a bill of rights for the internet, as sympathetic as I am to the idea. Rather, my point is that the most effective approach to protecting the right of privacy must be substantially international. Given the international nature of the net, a purely national response will be limited in its ability to contend with the subtle and complex challenges to privacy thrown up by our digital interactions.

However, the law is not the only regulatory modality. Social norms, market forces and coding all play a significant role in affecting behaviour, and hence in protecting (or not protecting) privacy. These too require a mix of international and national responses. Hence it has been suggested that codes of conduct, such as Brazil's Marco Civil da Internet, might help to articulate appropriate social norms. Such codes of conduct could well co-exist alongside bills of rights, the former as guidelines, the latter providing legally enforceable rights. Together, they would also help to shape appropriate market and coding initiatives. Whatever specific forms they might ultimately take, a code of conduct and a bill of rights would be significant positive steps. Young users could discuss their provisions; adult users could be reminded of their clauses in an iterative manner; and relevant provisions could be incorporated into national and international law. Kant's formula of humanity offers an excellent foundation. A code of conduct and a bill of rights formulated on the dual principles of not using others merely as means, and of respecting others as end-setters, would provide a powerful booster of privacy, among other freedoms. This would help users to assert their rights and

responsibilities not just as citizens of individual countries, but as “netizens”: the citizens of a globally connected internet (MacKinnon, 2012).

In the very optic fibres of its being, the internet wants cosmopolitanism. Only a significantly international approach can protect privacy in the context of a globe-spanning internet, and only with such protection might our world, or some part of it, begin to approach cosmopolitanism. Only thus can our goal of *cosmopolis* be joined by a goal of *cosmoikos*. We might dub this synthesis of goals *cosmoikopolis*, a blend of *cosmos* (world), *oikos* (private) and *polis* (public). In *cosmoikopolis*, private and public are valued globally and in balance. I have already said that Kant was an optimist about the future of humanity and morality. Central to Kant’s philosophy is the idea that people are capable of moral improvement, and that on a larger scale, humanity has the potential for moral progress (see Anderson-Gold, 2010: 32). In the face of the significant challenges posed by our online interactions, our goal ought to be *cosmoikopolis*, where citizens globally enjoy a relational privacy that is universally valued, even as it varies by region.

Conclusion

In Panopticon 2.0, privacy is under threat. The *condition* of privacy has suffered, and continues to suffer. However, the *right* to privacy persists. Indeed, the right to privacy appears to be strengthening, protected by an array of instruments ranging from revenge porn laws to Brazil’s Marco Civil da Internet to the sweeping provisions of Europe’s General Data Protection Regulation. This right to privacy, I have been arguing, ought to be further bolstered by legal protections modeled on consumer law, and also by the creation of a tort of serious invasion of privacy (in jurisdictions where none exists). However, legal responses must be supplemented by extra-legal responses. Following Lessig, I have argued that the law is a significant regulator of the internet, but that supplementary regulators exist in the shape of social norms, market forces and coding. A two-tier model of consent, which includes individual consent and the collective consent of the law, enables us to enforce Kant’s mere means principle; but it is often through extra-legal measures that Kant’s ends in themselves principle can be fostered. It is partly

through clearly-articulated privacy guidelines that promote respect for one another, for instance, that a due respect for privacy might come to prevail. Challenges to privacy are also challenges to our relationships, and legal remedies alone are insufficient. Drawing on the formula of humanity, consent was my starting point. Certainly, there exists ample scope for coding consent into the internet, and social norms and market forces can empower consent too. But consent isn't enough. There are other values at stake than those which can be protected by consent, including friendship. Hence these four regulatory modalities begin with consent, but do not end there.

In *The Social Network* (as in real life), Mark Zuckerberg is aided in his efforts to develop Facebook by Sean Parker, who had previously run the file-sharing music service Napster. "We lived on farms, then we lived in cities, and now we're going to live on the internet," Parker says. This, he says, has significant privacy ramifications. "Whatever it is that's gonna trip you up, you've done already. Private behaviour is a relic of a time gone by." The dialogue is fictional, but the "privacy is dead" sentiment is widely-held. In response, I have been arguing that privacy *can* be protected, alongside other freedoms. The challenges posed by our online interactions threaten privacy, but are also accompanied by unprecedented opportunities. In its potential for a global connectivity, in its capacity for all humanity to engage with one another on an equal footing, in its promise for each person to treat every other as a fellow user with all contact constructed entirely and equitably of ones and zeros, the internet is a place well suited to approximating Kant's ideal of a cosmopolitan realm of ends. A place where we can treat one another, and ourselves, as *persons*.

Conclusion

Each day, the online shop tmart.com sends me an email with its latest offerings of gadgets and gizmos. On November 3, 2016, the subject line of the email promised, “Mini Spy Clothes Hook Hidden Camera.” For \$17.99, beneath offers for a flashlight and walkie talkie, the email offered an oval-shaped piece of plastic with two hooks and a tiny indent concealing a hidden camera. What was its purpose? Catching shoplifters? Filming people in change rooms? Here was just one more tiny reminder of how the internet and digital technology are challenging privacy, and how, more than ever, privacy norms need to be articulated and observed.

Three main questions motivated this thesis. First, how does the internet confuse and challenge privacy? Second, what is privacy, and why should we care? And third, is there a normative principle we can use to better understand and protect privacy on the internet? These three questions were addressed, respectively, in chapters one and two, chapters three and four and chapters five and six.

In chapter one, I described how our internet interactions are serving to confuse and challenge privacy. In part, this is because the internet is enabling a multiplication of place, in which users can simultaneously find themselves in several spaces, some physical and some virtual, all layered on top of one another. The confusion and challenge can arise when public and private spaces collide, such as when I post to Twitter from my bedroom. Further, I argued that the internet is tending to convergence, ubiquity and multi-directionality, and that these three characteristics are tending to make data public. The result, I showed, is that surveillance has been joined by sousveillance and lateral viewing. This viewing extends into the past, and even has predictive power. Everyone can potentially watch everyone. I dubbed this phenomenon Panopticon 2.0.

However, any ensuing omniscience, in which everyone can know everything about everyone, is merely potential. In reality there are limits to watching, which I detailed in chapter two. These include guerrilla tactics adopted by net users, such as strategies of obfuscation or using browsers such as TOR. They also involve a

recognition that the challenge to privacy arises from three sources: from individuals; from companies; and from governments. Each of these presents distinctive challenges and requires distinctive responses. In the second section of chapter two, I showed that the internet is the site of an epic clash of norms and values. On the net, user-generated “norms” vary wildly, often accompanied by the claim that real world norms ought not apply. Meanwhile, a series of values have been coded into the architecture of the internet and its platforms, and these have tended to promote openness, not privacy. Given the cumulative effect of these user-generated norms and embedded values, I proposed that we urgently need to articulate which privacy norms and values ought to apply on the net.

In chapter three, I turned to the daunting task of defining privacy. The term privacy has come to attach itself to an increasingly wide range of information, situations, activities and personal attributes. From the ancient Greek distinction between *oikos* and *polis*, in which privacy pertained to place, the term has grown to yield a *right* to privacy, which offers a legal and ethical entitlement encompassing individual privacy. Seeking a conceptual account, I explored several models to argue that privacy can be defined by *restrictions on access*, which are sometimes about control. In chapter four, I argued that privacy matters, for at least three reasons. The first is dignity. If a person is being spied upon without their knowledge, I argued, there is a harm being committed, and that person’s dignity has not been respected. The second is autonomy. In many cases, a *non-consensual* invasion of privacy will lead people to act or think differently. The third lies in our relationships. Without due privacy, our ability to love, trust and befriend is compromised. Further, I proposed the concept of *relational* privacy, which recognises that each of us is socially embedded, and that privacy does not merely allow for separation and withdrawal from others, but also helps to inform and solidify our connections to others.

Having shown that the internet challenges privacy, that privacy is about limitations upon access and that privacy is relational and matters deeply, I turned to Kant’s formula of humanity. After demonstrating its suitability as a normative principle for internet privacy, I argued that the formula’s application relies on individual consent. Whenever consent is required, the first issue involves competence. In cases of competence, I proposed and developed a model of actual

consent, so defined as also to incorporate elements of possible consent. Even so, individual consent requires oversight by the collective consent of the law. Given that data flows so easily and unpredictably online, collective consent is required in the shape of *just* laws to enact the right to privacy. This two-tier model of consent aligns with the access model of privacy adopted in chapter three: individual consent involves control over access; collective consent involves externally-imposed limitations upon access.

In chapter six, I proposed a series of forms that just laws ought to take, including the requirement that a national approach be supplemented with an international approach. After returning to the triple threat from individuals, companies and governments, I argued that a tort of serious invasions of privacy would be sufficiently flexible to respond to these various threats, and should be enacted in jurisdictions where no such remedy currently exists. I also argued for privacy protections modeled on consumer protections, which offer both civil and criminal remedies, as appropriate. To give substance to these protections based on consumer law, I proposed a series of privacy principles that promote consent, transparency and fairness, that outlaw coercion and deception, and that allow that privacy must be balanced against other rights and interests. However, I argued that not every privacy issue ought involve the law, and that not every privacy issue ought involve consent. Hence I suggested a number of supplementary measures to protect internet privacy, within the extra-legal modalities of social norms, market forces and coding. With the formula of humanity as a base and with a response that is appropriately international, privacy norms can be articulated, fostered and enforced in a way that is universal, and yet that allows for significant variation from culture to culture. In this way, the due respect of privacy can bring us nearer Kant's *cosmopolis*.

A provisional title for this thesis was, "Why new media needs old ethics." In some ways, however, invoking Kant felt counter-intuitive. As I noted in chapter one, the internet is marked out by a relentless pace of change. Its protean nature is potentially problematic for ethicists. As Shannon Vallor writes:

The founders of the most enduring classical traditions of ethics – Plato, Aristotle, Aquinas, Confucius, the Buddha – had the luxury of assuming that the practical

conditions under which they and their cohorts lived would be, if not wholly static, at least relatively stable (Vallor, 2016: 6).

On the internet, flux seems to be the only constant. Still, I hope to have shown that new media *does* need old ethics. Or at least that venerable ethical principles can illuminate a context far beyond any imaginable for their authors. And perhaps, by extension, this thesis might even suggest that our online interactions generally, and not just those that relate to privacy, might benefit from a broader code of conduct, based on principles that include (and are perhaps founded on) the formula of humanity.

Of course, there is much work to be done. First, the concept of relational privacy requires significant clarification and elaboration. Second, the formula of humanity is not satisfied by the application of consent and the mere means principle. The formula's supplementary prescription, the ends in themselves principle, involves the application of love, respect and self-respect. One particularly rich prospect for investigation, which I have signaled but not explored, involves applying to the internet the argument that privacy matters for our social relationships. For instance, just how is Facebook harming our friendships, and our ability to befriend? Third, this theoretical work might then inform a more detailed account of potential practical responses. I have suggested legal and extra-legal measures, yet my recommendations are anything but exhaustive. The formula of humanity can yield further proposals. Fourth, the recommendations I have made require fleshing out. Most urgently, perhaps, legal drafting is needed to devise privacy protections modeled on consumer law. Fifth, a coherent and more detailed international response ought to be developed, both on a legal front, and an extra-legal front, to begin to enact effective *global* protections of privacy.

Throughout this thesis, I have invoked novels and films that illuminate my topic, teasing out various issues for closer inspection. The anthropomorphised operating system of *Her* suggested how networked technology could impact privacy (and our relationships) in profound and subtle ways. The predictive policing of *Minority Report* prompted me to argue that ethical questions need to be addressed *before* new technologies are implemented. *Rear Window* and *The Truman Show* showed that privacy is about access, and that it matters a great deal. *The Dark Knight*, as a metaphor for blanket surveillance by government agencies, exposed

the connection between individual consent and collective consent; and *The Social Network* showed that privacy is about consent, but not *only* about consent. Meanwhile, debate rages in the literature. With its utopian vision of Gaia, *Foundation's Edge* paints privacy as an obstacle for humanity; whereas *1984* and *The Knife of Never Letting Go*, with its account of the “Noise”, both depict worlds without privacy as dystopian. The internet has tremendous, unprecedented potential. My arguments in favour of privacy should not be taken as arguments against the internet. My argument is simply that we must work to create a future that, with all its networked connectivity, does not begin to resemble the “Noise” or *1984*. To do so, the formula of humanity can help.

In this thesis, I have mounted a defence of privacy, and of the right to privacy. Perhaps, however, we need to think less of the right, and more of the *duty* of privacy. Rights, as we know, come with concomitant duties. There cannot be rights without respective obligations. Once we start thinking of a duty to respect privacy, a duty that we share individually and collectively, then we might more easily recognise the value of privacy. As I have shown, the way we fulfill that duty of respecting individual privacy may permissibly vary from person to person, from country to country, but the observance of at least one fundamental principle ought to remain constant. In line with the formula of humanity, we must treat one another always as end-setting agents, and never merely as means. Hence the use of certain hidden cameras, such as those concealed in clothes hooks, ought to be limited, and perhaps even outlawed. On these occasions, the formula of humanity can help to ensure that we treat one another more ethically, by observing the duty to respect one another's privacy.

Appendix

The preparation of this thesis involved primary research in the form of five interviews conducted during 2015 and 2016. These interviews were with: Timothy Pilgrim, Australia's privacy commissioner; Samantha Yorke, public policy and government relations counsel at Google Australia; Benjamin Carr, chief privacy officer at Telstra; Nigel Waters, committee member of the Australian Privacy Foundation; and a spokesperson for the Federal Attorney-General's Department. The ethical aspects of this research were approved by the Macquarie University Human Research Ethics Committee, reference number 5201500524.

The aim of the interviews was to develop and test ideas raised by the research and also to reveal how various stakeholders think about privacy. The interviewees were deliberately selected to provide a range of perspectives aligning with the triple threat to privacy identified in chapters two and six: from individuals (the Australian Privacy Commissioner, a privacy advocate); from companies (Google, Telstra); and from governments (the Australian Federal Attorney-General's Department). More specifically, the interview questions were designed to test the theoretical framework developed in chapter five, which involves overlaying individual consent with the collective consent of legal regulation, and to elicit responses on various practical solutions proposed in chapter six, including the potential role to be played by privacy protections modeled on consumer protections. Interviewees were also asked further questions as to the meaning and value of privacy. Further interviews were sought, but not obtained, with Facebook, Edward Snowden and Julian Assange.

In the thesis, the interviews are not printed in full. Rather, they are used to clarify issues as they arise. Quotes from the interviews appear mainly in chapter six, which involves applying the normative framework developed earlier and hence proposing practical solutions. These interviews were also intended to draw on my experience as a journalist.

In each case, contact was initially made by email during the second half of 2015. Two interviews were conducted face-to-face; two were conducted by email; and

one was conducted by telephone. The face-to-face and phone interviews were recorded on a digital recording device (a Zoom recorder) and later transcribed onto my personal computer.

Three interviews were conducted as follows:

- *Face-to-face interview with Timothy Pilgrim, Australian Privacy Commissioner, Thursday, October 8, 2015*
- *Face-to-face interview with Benjamin Carr, chief privacy officer, Telstra, Wednesday, December 9, 2015*
- *Telephone interview with Nigel Waters, committee member of the Australian Privacy Foundation, Thursday, December 10, 2015*

Two sets of emailed responses were received as follows:

- *Email responses received from Samantha Yorke, public policy and government relations counsel, Google Australia, Tuesday, November 3, 2015*
- *Email responses received from a spokesperson for the Attorney-General's Department in Canberra, Thursday, March 31, 2016*

The consent form, list of indicative questions and approval form from the Macquarie University Human Research Ethics Committee can be found overleaf.

DEPARTMENT OF PHILOSOPHY

Faculty of Arts
Macquarie University
NSW 2109 Australia
T: +61 (2) 9850 8837
F: +61 (2) 9850 8892
www.phil.mq.edu.au
ABN 90 952 801 237
CRICOS Provider No 00002J



MACQUARIE
University
SYDNEY · AUSTRALIA

Participant Information and Consent Form

Project Title: The Ethics of Internet Privacy: Applying Kant's Formula of Humanity to our Online Interactions

Research Aims: The internet is changing the way people interact and challenging ethical norms, including norms relating to privacy. My thesis aims to interrogate the meaning, the value and the protection of internet privacy. In particular, I aim to explore the extent to which principles of consent and the public interest can illuminate the ethics of internet privacy. The aim of these interviews, which will draw on my extensive experience as a journalist at the Sydney Morning Herald, is to gain further insight into these issues by talking to people with significant expertise in areas related to internet privacy. The research will thus involve interviews with key figures from the media industry and beyond.

Researcher: Sacha Molitorisz
Philosophy Department, Macquarie University
0423 306 769
sm4860@nyu.edu

This research is being conducted to meet the requirements for the degree of PhD under the supervision of Professor Catriona Mackenzie (Department of Philosophy, 9850 8865, catriona.mackenzie@mq.edu.au), Dr Paul Formosa (Department of Philosophy, 9850 8817, paul.formosa@mq.edu.au) and Professor Sherman Young (Department of Media, Music, Communication and Cultural Studies, 9850 6778, sherman.young@mq.edu.au).

Participants who agree to be involved will be required to be interviewed for approximately 45 minutes on the subject of internet privacy. The interviews will be face-to-face, if possible, at a venue convenient for the participant. If face-to-face is impossible, the interviews can be conducted via phone or email. Interviews done face-to-face or via email will be recorded on a Zoom audio recorder.

Portions of the interviews will be published in the completed PhD. They may also be used for conference papers and presentations, and also if the PhD is developed into a book. Participants have the option of being quoted by name (the preferred option), or of being quoted anonymously.

This research is funded by an Australian Postgraduate Award.

Participation in this research is entirely voluntary. Participants can withdraw from the process at any time without adverse consequence.

[Please see next page]

I, _____, agree to participate in this research.

Date:

SIGNED:

Interviewee/participant:

Signature:

Interviewer/researcher: Sacha Molitorisz

Signature:

Two copies of this document have been signed. The researcher and the interviewee will each keep one copy.

Note: The ethical aspects of this study have been approved by the Macquarie University Human Research Ethics Committee. If you have any complaints or reservations about any ethical aspect of your participation in this research, you may contact the Committee through the Director, Research Ethics and Integrity (phone 9850 7850; email ethics@mq.edu.au). Any complaint you make will be treated in confidence and investigated, and you will be informed of the outcome.

Version 1, Wednesday, August 12, 2015

Interview Questions

1. How do you/does your organisation define privacy?
2. Why and how much does privacy matter to you/your organisation?
3. How do you/does your organisation protect the privacy of individuals?
4. Do you/Does your organisation think the internet has an architecture of openness?
5. Do you/Does your corporation believe privacy norms are shifting and that people are becoming more willing to share more personal information?
6. How do you/does your organisation think privacy ought to be protected, if at all?
7. When it comes to internet privacy, what role do you/does your organisation see for individual consent?
8. When it comes to internet privacy, what role do you/does your organisation see for regulation and the law?
9. Do you/does your organisation think that if people compromise their privacy, they may be compromising their ability to trust, to befriend, to love and to form relationships?
10. Should people have an ethical and/or legal right to privacy?
11. Should there be laws/regulations to prevent individuals/companies/governments from being misleading, deceptive or coercive when dealing with matters pertaining to people's privacy?

If necessary, these questions will be supplemented with follow-up questions that pursue the lines of enquiry outlined above, probing the meaning, value and protection of privacy. This will be in order to clarify the issues raised above.

Office of the Deputy Vice-Chancellor (Research)

Research Office
Research Hub, Building C5C East
Macquarie University
NSW 2109 Australia
T: +61 (2) 9850 7987
<http://www.research.mq.edu.au/>
ABN 90 952 801 237
CRICOS Provider No 00002J



MACQUARIE
University
SYDNEY · AUSTRALIA

12 August 2015

Professor Catriona Mackenzie
Macquarie University NSW 2109

Dear Professor Mackenzie,

Reference No: 5201500524

Title: *The Ethics of Internet Privacy: Applying Kant's Formula of Humanity to our Online Interactions*

Thank you for submitting the above application for ethical and scientific review. Your application was considered by the Macquarie University Human Research Ethics Committee (HREC (Human Sciences & Humanities)) at its meeting on 31 July 2015 at which further information was requested to be reviewed by the Ethics Secretariat.

The requested information was received with correspondence on 12 August 2015.

I am pleased to advise that ethical and scientific approval has been granted for this project to be conducted at:

- Macquarie University

This research meets the requirements set out in the *National Statement on Ethical Conduct in Human Research* (2007 – Updated March 2014) (the *National Statement*).

This letter constitutes ethical and scientific approval only.

Standard Conditions of Approval:

1. Continuing compliance with the requirements of the *National Statement*, which is available at the following website:

<http://www.nhmrc.gov.au/book/national-statement-ethical-conduct-human-research>

2. This approval is valid for five (5) years, subject to the submission of annual reports. Please submit your reports on the anniversary of the approval for this protocol.

3. All adverse events, including events which might affect the continued ethical and scientific acceptability of the project, must be reported to the HREC within 72 hours.

4. Proposed changes to the protocol must be submitted to the Committee for approval before implementation.

It is the responsibility of the Chief investigator to retain a copy of all documentation related to this project and to forward a copy of this approval letter to all personnel listed on the project.

Should you have any queries regarding your project, please contact the Ethics Secretariat on 9850 4194 or by email ethics.secretariat@mq.edu.au

The HREC (Human Sciences and Humanities) Terms of Reference and Standard Operating Procedures are available from the Research Office website at:

http://www.research.mq.edu.au/for/researchers/how_to_obtain_ethics_approval/human_research_ethics

The HREC (Human Sciences and Humanities) wishes you every success in your research.

Yours sincerely

Dr Karolyn White

Director, Research Ethics & Integrity,
Chair, Human Research Ethics Committee (Human Sciences and Humanities)

This HREC is constituted and operates in accordance with the National Health and Medical Research Council's (NHMRC) *National Statement on Ethical Conduct in Human Research* (2007) and the *CPMP/ICH Note for Guidance on Good Clinical Practice*.

Details of this approval are as follows:

Approval Date: 12 August 2015

The following documentation has been reviewed and approved by the HREC (Human Sciences & Humanities):

Documents reviewed	Version no.	Date
Macquarie University Ethics Application Form	2.3	Received 24/06/2015
Correspondence from Ms Sacha Molitorisz responding to the issues raised by the HREC (Human Sciences and Humanities)		Received 12/08/2015
MQ Participant Information and Consent Form (PICF)	1	12/08/2015
Interview Questions	1	12/08/2015

Bibliography

- 4chan. (2016), 4chan.org. (accessed 21 July 2016, 1.01pm).
- @CreepShot. (2017) @CreepShot, twitpic.com. (accessed 15 March 2017, 6.28pm).
- A.L.C. (1920) "When silence gives consent". *The Yale Law Journal* 29: 441-444.
- AAP. (2016a) "A California judge has ruled a lawsuit over Facebook's collection of biometric data from people 'tagged' in photos posted by other users can proceed". sbs.com.au. (accessed 11 August 2016, 10.13pm).
- AAP. (2016b) "France has fined Google over its refusal to delist results from its non-European websites such as Google.com". 25 March 2017, sbs.com.au. (accessed 21 February 2017, 4.14pm).
- Abokhodair N, Abbar S, Vieweg S, et al. (2016) "Privacy and twitter in Qatar: traditional values in the digital world. *Proceedings of the 8th ACM Conference on Web Science*. ACM, 66-77.
- Acquisti A and Fong CM. (2015) "An experiment in hiring discrimination via online social networks". 17 July 2015, ssrn.com/abstract=2031979. (accessed 17 August 2016, 3.30pm).
- Acquisti A, Gross R and Stutzman F. (2014) Face recognition and privacy in the age of augmented reality. *Journal of Privacy and Confidentiality* 6: 1-20.
- Axciom. (2016) "Axciom Privacy Principles". axciom.com. (accessed 16 September 2016, 2.31pm).
- Albrechtslund A. (2008) "Online social networking as participatory surveillance". *First Monday*, 13, firstmonday.org. (accessed 5 August 2016, 1.28pm).
- ALII. (2016a) "NSW Consolidated Acts: *Crimes Act 1900 - sect. 61N*". Australasian Legal Information Institute, austlii.edu.au. (accessed 9 February 2017, 12.17pm).
- ALII. (2016b) "NSW Consolidated Acts: *Inclosed Lands Protection Act 1901 - sect. 4*". Australasian Legal Information Institute, austlii.edu.au. (accessed 12 September 2016, 5.42pm).
- ALII. (2017a) "Commonwealth Consolidated Acts: *Competition and Consumer Act 2010 - Schedule 2*". Australasian Legal Information Institute, austlii.edu.au. (accessed 9 February 2017, 12.59pm).

- ALII. (2017b) "Commonwealth Consolidated Acts: *Criminal Code Act 1995* - Schedule". Australasian Legal Information Institute, austlii.edu.au.
- AlJazeera. (2015) "Can Kuwait justify mandatory DNA testing?". *Al Jazeera*, 14 July 2015. (accessed 23 November 2016, 4.55pm).
- Allen AL. (1988) *Uneasy access: privacy for women in a free society*, Lanham, US: Rowman & Littlefield.
- Allison HE. (2011) *Kant's Groundwork for the metaphysics of morals: a commentary*, Oxford, UK: Oxford University Press.
- Almuhimedi H, Schaub F, Sadeh N, et al. (2015) "Your location has been shared 5,398 times! A field study on mobile app privacy nudging". *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*. ACM, 787-796.
- ALRC. (2003) Essentially yours: the protection of human genetic information in Australia, 30 May 2003. Australian Law Reform Commission.
- ALRC. (2007) Review of Australian privacy law. Australian Law Reform Commission.
- ALRC. (2014) *Serious invasions of privacy in the digital era*. Australian Law Reform Commission, 31 March 2014, alrc.gov.au. (accessed 6 April 2017, 7.32pm).
- Anderson-Gold S. (2010) "Privacy, respect and the virtues of reticence in Kant". *Kantian Review* 15: 28-42.
- Anitha S and Gill A. (2009) "Coercion, consent and the forced marriage debate in the UK". *Feminist legal studies* 17: 165-184.
- Anthony S. (2016) "Facebook wins privacy case, can track any Belgian it wants". *Ars Technica*, 30 June 2016. (accessed 11 August 2016, 10.21pm).
- Antón AI, He Q and Baumer DL. (2004) "Inside JetBlue's privacy policy violations". *IEEE Security and Privacy* 2: 12-18.
- APS. (2015) *Stress and wellbeing: how Australians are coping with life*. Australian Psychological Society, psychology.org.au. (accessed 2 May 2017, 8.14pm).
- Arendt H. (1977) "Public rights and private interests". In: Mooney M and Stuber F (eds) *Small comforts for hard times: Humanists on public policy*. New York: Columbia University Press, 103-108.

- Arseniev-Koehler A, Lee H, McCormick T, et al. (2016) "#Proana: pro-eating disorder socialization on Twitter". *Journal of Adolescent Health* 58: 659-664.
- Asimov I. (1982) *Foundation's Edge*, London, UK: Granada.
- Assange J. (2013) "The banality of 'Don't Be Evil' ". *The New York Times*, 1 June 2013. (accessed 18 October 2016, 2.37pm).
- Attorney-General. (2016) Email interview between Sacha Molitorisz and a spokesperson for the Attorney-General's Department in Canberra, 31 March 2016.
- Audi R. (ed., 1999) *The Cambridge Dictionary of Philosophy*, Cambridge, UK: Cambridge University Press.
- Augoustinos M. (1999) "Ideology, false consciousness and psychology". *Theory & Psychology* 9: 295-312.
- AustralianGovernment. (2017) *Telecommunications (Interception and Access) Amendment (Data Retention) Act 2015*. Federal Register of Legislation, legislation.gov.au. (accessed 5 April 2017, 10.04pm).
- Aveyard H. (2002) "Implied consent prior to nursing care procedures". *Journal of Advanced Nursing* 39: 201-207.
- Baek TH and Morimoto M. (2012) "Stay away from me". *Journal of advertising* 41: 59-76.
- Baker R. (2015) "Facebook makes data deal with Quantum, Acxiom and Experian to fuse offline and online data". *AdNews*, 21 July 2015. (accessed 12 August 2016, 5.57pm).
- Ball J. (2014) "GCHQ has tools to manipulate online information, leaked documents show". *The Guardian*, 15 July 2014. (accessed 6 February 2017, 4.57pm).
- Bambauer DE. (2013) "Privacy versus security". *Journal of Criminal Law and Criminology* 103: 667-684.
- Bane AF and Milheim WD. (1995) "Internet insights: how academics are using the Internet". *Computers in libraries* 15: 32-36.
- Barker E. (2016) "This is the best way to overcome Fear Of Missing Out". *Time*, 7 June 2016.
- Barrett D. (2015) "What is the law on revenge porn?". *Telegraph.co.uk*. (accessed 2 February 2016, 1.53pm).

- BBC. (2005) "Classic Kiss sold at auction". *BBC News*, 25 April 2005. (accessed 2 November 2016, 4.25pm).
- BBC. (2015) "How I was 'cyber-flashed' ". *BBC News*, 13 August 2015. (accessed 18 October 2016, 2.28pm).
- Beauchamp TL and Childress JF. (2001) *Principles of biomedical ethics*, New York, US: Oxford University Press.
- Behar R. (2004) "Never Heard Of Acxiom? Chances Are It's Heard Of You". *Fortune - European Edition*, 149. (accessed 11 August 2016, 2.21pm).
- Behzadan AH, Dong S and Kamat VR. (2016) "Recent advances in augmented reality for architecture, engineering, and construction applications". In: Barfield W (ed) *Fundamentals of wearable computers and augmented reality*. 2 ed. Boca Raton, US: CRC Press.
- Benn SI. (1971) Privacy, freedom, and respect for persons. In: Pennock JR and Chapman JW (eds) *Nomos XIII: Privacy*. New York, US: Atherton Press, 1-26.
- Bennett CJ. (1992) *Regulating privacy: data protection and public policy in Europe and the United States*, Ithaca, NY, US: Cornell University Press.
- Bentham J. (1811) "Outline of the plan of construction, alluded to in the 'Proposal for a new and less expensive mode of employing and reforming convicts'". *Remarks on the Form and Construction of Prisons: With Appropriate Designs*. London, UK: The Committee of the Society for the Improvement of Prison Discipline etc.
- Berger A. (2009) "Brandeis and the history of transparency". *Sunlight Research*, 26 May 2009. (accessed 21 September 2016, 8.12pm).
- Bergman MK. (2001) "White paper: the deep web: surfacing hidden value". *Journal of electronic publishing*, 7: 1, [dx.doi.org/10.3998/3336451.0007.104](https://doi.org/10.3998/3336451.0007.104). (accessed 15 May 2017, 11.07am).
- Berners-Lee T. (1999) *Weaving the Web: The original design and ultimate destiny of the World Wide Web by its inventor*, San Francisco, US: HarperInformation.
- Berners-Lee T, Hendler J and Lassila O. (2001) "The semantic web". *Scientific american* 284: 28-37.
- Bloustein EJ. (1964) "Privacy as an aspect of human dignity: an answer to Dean Prosser". *New York University Law Review* 39: 962-1007.

- Blue V. (2013) "Firm: Facebook's shadow profiles are 'frightening' dossiers on everyone". *ZDNet*, 24 June 2013. (accessed 11 August 2016, 6.49pm).
- Bok S. (1982) *Secrets: on the ethics of concealment and revelation*, New York, US: Pantheon.
- Bosker B. (2013) "'Hell is other people' helps you steer clear of your awful 'friends'". *The Huffington Post*, 19 June 2013. (accessed 23 September 2016, 3.16pm).
- boyd d. (2014) It's complicated: the social lives of networked teens. *danah.org*. (accessed 20 April 2017, 8.19pm).
- Brand M. (2016) "Can Facebook influence an election result?". *The Conversation*, 27 September 2016. (accessed 28 September 2016, 1.30pm).
- Brand M, Laier C and Young KS. (2014) "Internet addiction: coping styles, expectancies, and treatment implications". *Frontiers in psychology* 5: 1256.
- Brandeis LD. (2009) *Other people's money and how the bankers use it*, New York, US: Cosimo, Inc.
- Brandimarte L, Acquisti A and Loewenstein G. (2013) "Misplaced confidences: privacy and the control paradox". *Social Psychological and Personality Science* 4: 340-347.
- Brin D. (1998) *The transparent society: Will technology force us to choose between privacy and freedom?*, Reading, Mass., US: Perseus Books.
- Brown W. (2004) "'The Subject of Privacy': A Comment on Moira Gatens". In: Rössler B (ed) *Privacies: Philosophical Evaluations*. Stanford, US: Stanford University Press, 133-141.
- Bruner R. (2016) "A brief history of Instagram's fateful first day". *Time*, 16 July 2016. (accessed 1 February 2017, 6.42pm).
- Bruns A. (2007) "Produsage: towards a broader framework for user-led content creation". *Proceedings of the 6th ACM SIGCHI conference on Creativity & cognition*. ACM, 99-106.
- Brustein J. (2015) "RadioShack's bankruptcy could give your customer data to the highest bidder". *Bloomberg*, 25 March 2015. (accessed 17 November 2016, 3.22pm).
- Calkins MM. (2000) "They shoot Trojan horses, don't they? An economic analysis of anti-hacking regulatory models". *Georgetown Law Journal* 89: 171-224.

- Calpito D. (2015) "Google received 350,000 right to be forgotten requests for 1.2 million links in Europe". *Tech Times*, 26 November 2015. (accessed 1 December 2016, 1.39pm).
- Canning S. (2016) "Facebook and Quantum sign deal to measure advertising impact on store sales". *Mumbrella*, 8 June 2016. (accessed 15 August 2016, 5.24pm).
- Carr B. (2015) Face-to-face interview with Sacha Molitorisz, 9 December 2015.
- Castells M. (1996) *The rise of the network society: The information age: Economy, society, and culture*, Oxford, UK: John Wiley & Sons.
- CDT. (2008) "Existing federal privacy laws". *Center for Democracy and Technology*, 30 November 2008, cdt.org. (accessed 14 February 2017, 10.24am).
- Chadwick R and Berg K. (2001) "Solidarity and equity: new ethical frameworks for genetic databases". *Nature Reviews Genetics* 2: 318-321.
- Chatzimilioudis G, Konstantinidis A, Laoudias C, et al. (2012) "Crowdsourcing with smartphones". *IEEE Internet Computing* 16: 36-44.
- Chideya F. (2015) "Your data is showing: breaches wreak havoc while the government plays catch-up". *The Intercept*, 28 May 2015. (accessed 9 August 2016, 6.33pm).
- Chighine L. (2015) "Wilson v Ferguson [2015] WASC 15 - remedy of equitable compensation for breach of confidence where the damage suffered is embarrassment, anxiety and distress". *Intellectual Property Forum: Journal of the Intellectual and Industrial Property Society of Australia and New Zealand*. 64-67.
- Christians CG, Fackler M, Richardson K, et al. (2012) *Media ethics: Cases and moral reasoning*, Glenview, Ill., US: Pearson Education.
- Citizenfour. (2014) Dir: Laura Poitras. US/Germany: Radius-TWC, 114mins.
- Cocking D and Kennett J. (1998) "Friendship and the self". *Ethics* 108: 502-527.
- Coleman G. (2014) *Hacker, hoaxer, whistleblower, spy: the many faces of Anonymous*, London, UK; New York, US: Verso Books.
- CommonSenseMedia. (2009) "Digital literacy and citizenship in the 21st century: educating, empowering and protecting America's kids". *A Common Sense Media White Paper*, itu.int, June 2009. (accessed 10 February 2017, 6.35pm).

- Cooke R. (2013) "How NSA surveillance destroys privacy and undermines our sovereignty". *The Monthly*, July 2013. (accessed 19 August 2016, 11.11am).
- Cormack L. (2017) "Choice calls for crackdown on unreasonably lengthy online consumer contracts". *smh.com.au*, 15 March 2017. (accessed 15 March 2017, 12.23pm).
- Corrigan R. (2015) "Mass surveillance not effective for finding terrorists". *New Scientist*, 15 January 2015, newscientist.com. (accessed 11 April 2017, 12.44pm).
- Cranor LF and Reidenberg J. (2002) "Can user agents accurately represent privacy notices?". *The 30th Research Conference on Communication, Information and Internet Policy (TPRC 2002)*. Alexandria, Virginia.
- Crawford A and Goggin G. (2009) "Geomobile web: Locative technologies and mobile media". *Australian Journal of Communication* 36: 97-109.
- CreepShots. (2017) "About: Photographer Rights". *creepshots.com*. (accessed 16 March 2017, 3.02pm).
- CriticalCommons. (2016) "Dark Knight cell phone surveillance". Dark Knight Montage, *criticalcommons.org*. (accessed 19 October 2016, 6.05pm).
- Curran G and Gibson M. (2013) "WikiLeaks, anarchism and technologies of dissent". *Antipode* 45: 294-314.
- Daniel C and Palmer M. (2007) "Google's goal: to organise your daily life". *Financial Times*, 23 May 2007. (accessed 17 August 2016, 3.41pm).
- de Mars S and O'Callaghan P. (2016) "Privacy and search engines: forgetting or contextualizing?". *Journal of Law and Society* 43: 257-284.
- Dean A. (2016) "Age of consent laws". *Australian Institute of Family Studies*, April. (accessed 24 October 2016, 5.34pm).
- Dearden L. (2016) "Burkini ban: why is France arresting Muslim women for wearing full-body swimwear and why are people so angry?". *independent.co.uk*, 25 August 2016. (accessed 15 September 2016, 2.30pm).
- Debatin B, Lovejoy JP, Horn AK, et al. (2009) "Facebook and online privacy: Attitudes, behaviors, and unintended consequences". *Journal of Computer - Mediated Communication* 15: 83-108.
- DeCew J. (2015) "Privacy". In: Zalta EN (ed) *The Stanford Encyclopedia of Philosophy*. plato.stanford.edu (Spring 2015 edition).

- Degeling M and Herrmann T. (2016) "Your interests according to Google: a profile-centered analysis for obfuscation of online tracking profiles". *Computers and Society*, arxiv.org. (accessed 21 February 2017, 4.38pm).
- Denis L. (2012) *Moral self-regard: Duties to oneself in Kant's moral theory*, New York, US: Routledge.
- DigitalLiteracy. (2017) "DigitalLiteracy.gov: your destination for digital literacy resources and collaboration". digitalliteracy.gov. (accessed 11 April 2017, 6.15pm).
- Dinello D. (2005) *Technophobia! Science fiction visions of posthuman technology*, Austin, Texas, US: University of Texas Press.
- Disconnect. (2016), disconnect.me. (accessed 22 August 2016, 4.33pm).
- Dropbox. (2015) "Dropbox terms of service". 4 November 2015. (accessed 16 August 2016, 6.29pm).
- Duckett C. (2016) "61 agencies after warrantless access to Australian telecommunications metadata". *ZDNet*. (accessed 18 August 2016, 5.59pm).
- Duhigg C. (2012) "How companies learn your secrets". *The New York Times*, 16. (accessed 16 August 2016, 12.02pm).
- Dutton WH. (1996) *Information and communication technologies: Visions and realities*, Oxford, UK: Oxford University Press.
- Dwyer T. (2015a) *Convergent media and privacy*, London, UK: Palgrave Macmillan.
- Dwyer T. (2015b) "Evolving concepts of personal privacy: locative media in online mobile spaces". In: Wilken R and Goggin G (eds) *Locative Media*. New York, US: Routledge, 121-135.
- Edutopia. (2015) "Digital citizenship: resource roundup". 21 October 2015. Edutopia.org. (accessed 10 February 2017, 7.10pm).
- Edwards CR. (2016) "Snap Inc. does what Google couldn't - exemplifies privacy by design". *Digerati*, 28 September 2016. (accessed 29 September 2016, 2.38pm).
- Edwards J. (2012) "How Facebook is hunting down and deleting fake accounts". *Business Insider Australia*, 30 December 2012. (accessed 12 August 2016, 3.43pm).

- Elgesem D. (1996) "Privacy, respect for persons, and risk". In: Ess C (ed) *Philosophical perspectives on computer-mediated communication*. Albany, NY, US: State University of NY Press, 45-66.
- Ellery D. (2016) "Ugg boot business fined \$10,800 for ACT manufacture claim". *smh.com.au*, 3 May 2016. (accessed 9 February 2017, 1.34pm).
- Elmer G. (2013) "IPO 2.0: The panopticon goes public". *MediaTropes*, 4. (accessed 30 August 2016, 6.42pm).
- Elmer G. (2015) "Going public on social media". *Social Media + Society*, 1. (accessed 11 August 2016, 6.41pm).
- Ember S. (2016) "Gawker and Hulk Hogan reach \$31million settlement". *nytimes.com*, 2 November 2016. (accessed 9 February 2017, 4.56pm).
- Eshet-Alkalai Y. (2004) Digital literacy: A conceptual framework for survival skills in the digital era. *Journal of Educational Multimedia and Hypermedia* 13 (1): 93-106.
- eSmart. (2017) "eSmart Digital Licence: teaching smart, safe, responsible online behaviour". *esmart.org.au*. (accessed 10 February 2017, 6.50pm).
- Ess C. (2011) "Self, community, and ethics in digital mediatized worlds". In: Ess C and Thorseth M (eds) *Trust and virtual worlds: Contemporary perspectives*. Bern, Switzerland: Peter Lang, 3-30.
- Etzioni A. (2000) "Social norms: Internalization, persuasion, and history". *Law and Society Review* 34, No. 1: 157-178.
- Etzioni A. (2015) "A cyber age privacy doctrine: more coherent, less subjective, and operational". *Brooklyn Law Review* 80: 1263-1308.
- Evers C and Goggin G. (2012) "Mobiles, men and migration: Mobile communication and everyday multiculturalism in Australia". In: Fortunati L, Perterra R and Vincent J (eds) *Migrations, diaspora and information technology in global societies*. New York, US: Routledge, 78-90.
- Facebook. (2016a) "How can I adjust how ads are targeted to me based on my activity off of Facebook?". Help Centre. (accessed 30 August 2016, 6.08pm).
- Facebook. (2016b) "How tagging works". (accessed 12 August 2016, 5.10pm).
- Fairfield JA and Engel C. (2015) "Privacy as a public good". *Duke Law Journal* 65 (3): 385-421.
- Fine C. (2008) *A mind of its own: How your brain distorts and deceives*, New York, US: WW Norton & Company.

- Fleischer P. (2015) Fireside chat with Peter Fleischer, Google's global privacy counsel, 24 October 2015. *Amsterdam Privacy Conference*.
- Foot P. (1972) "Morality as a system of hypothetical imperatives". *The Philosophical Review* 81: 305-316.
- Formosa P. (2008) " 'All politics must bend its knee before right': Kant on the relation of morals to politics". *Social Theory and Practice* 34: 157-181.
- Formosa P. (2013a) "Evils, wrongs and dignity: how to test a theory of evil". *The Journal of Value Inquiry* 47: 235-253.
- Formosa P. (2013b) "Kant's conception of personal autonomy". *Journal of Social Philosophy* 44: 193-212.
- Formosa P. (2017, in press) *Kantian ethics, dignity and perfection*, Cambridge, UK: Cambridge University Press.
- Foucault M. (1977) *Discipline and punish: The birth of the prison*, New York, US: Vintage.
- Franks MA. (2015) "Drafting an effective 'revenge porn' law: a guide for legislators". 17 August 2015, papers.ssrn.com. (accessed 7 April 2017, 1.26pm).
- Fried C. (1968) "Privacy". *Yale Law Journal* 77, No. 3: pp. 475-493.
- FTC. (2017) "Bureau of consumer protection". *Federal Trade Commission*, *ftc.gov*. (accessed 9 February 2017, 12.18pm).
- Gair K. (2015) "Privacy concerns mount as drones take to the skies". *smh.com.au*, 12 December 2015. (accessed 13 September 2016, 11.01am).
- Gallagher R and Greenwald G. (2014) "How the NSA plans to infect 'millions' of computers with malware". *The Intercept*, 12. (accessed 17 August 2016, 10.40pm).
- Gander K. (2014) "Miss Teen USA webcam hacker Jared James Abrahams sentenced to 18 months in prison". *The Independent*, 18 March 2014. (accessed 22 February 2017, 4.29pm).
- Gatford S. (2015) "Revenge porn makes new law". *LexisNexis*, 29 September 2015. (accessed 13 September 2016, 11.12am).
- Gavison R. (1980) "Privacy and the limits of law". *The Yale Law Journal* 89: 421-471.
- Gerety T. (1977) "Redefining privacy". *Harvard Civil Rights - Civil Liberties Law Review* 12: 233.
- Gerstein RS. (1978) "Intimacy and privacy". *Ethics* 89: 76-81.

- Glancy DJ. (1979) "The invention of the right to privacy". *Arizona Law Review* 21: 1-39.
- Godfrey M. (2013) "Revenge porn 'spreading like wildfire' ". *The Australian*, 22 November 2013. (accessed 2 February 2017, 2.03pm).
- Godwin M. (2003) *Cyber rights: Defending free speech in the digital age*, Cambridge, Mass., US: MIT press.
- Goggin G and Crawford K. (2010) "Moveable types: Youth and the emergence of mobile social media in Australia". *Media Asia* 37: 224-231.
- Goode L. (2015) "Anonymous and the political ethos of hacktivism". *Popular Communication* 13: 74-86.
- GoogleUK. (2014) "Google terms of service". 30 April 2014. (accessed 16 August 2016, 6.26pm).
- Goos K, Friedewald M, Webster W, et al. (2015) "The co-evolution of surveillance technologies and surveillance practices". In: Wright D and Kreissl R (eds) *Surveillance in Europe*. New York, US: Routledge, 51-100.
- Gordon E and de Souza e Silva A. (2011) *Net locality: Why location matters in a networked world*: John Wiley & Sons.
- Gould J and Hirst R. (2017) "Show me the money". In *Competition*, 15 February 2017, incompetition.com.au. (accessed 9 May 2017, 1.20pm).
- Greenberg A. (2014) "These are the emails Snowden sent to first introduce his epic NSA leaks". *Wired*, 13 October 2014. (accessed 17 August 2016, 10.53pm).
- Greenleaf G. (2015) "Global tables of data privacy laws and bills". *133 Privacy Laws & Business International Report*, 18-28; *UNSW Law Research Paper No. 2015-28*. 4th ed., January 30 2015.
- Greenwald G. (2013) "Edward Snowden: NSA whistleblower answers reader questions". *The Guardian*, 18 June 2016. (accessed 22 August 2016, 2.55pm).
- Greenwald G. (2014) *No place to hide: Edward Snowden, the NSA, and the US surveillance state*, New York, US: Metropolitan Books.
- Greenwald G and MacAskill E. (2013) "NSA Prism program taps in to user data of Apple, Google and others". *The Guardian*, 8 June 2013. (accessed 6 February 2017, 5.15pm).

- Grimmelmann J. (2009) "Privacy as product safety". *Widener Law Journal* 19: 793-827.
- Grisso T and Appelbaum PS. (1998) *Assessing competence to consent to treatment: a guide for physicians and other health professionals*, New York, US: Oxford University Press.
- Gross R and Acquisti A. (2005) "Information revelation and privacy in online social networks". *Proceedings of the 2005 ACM workshop on privacy in the electronic society*. ACM, 71-80.
- Gubbi J, Buyya R, Marusic S, et al. (2013) "Internet of Things (IoT): A vision, architectural elements, and future directions". *Future Generation Computer Systems* 29: 1645-1660.
- Guo Y and Zhang C. (2015) "Legal risks and solutions to e-marketers' data mining". In: Guo Y (ed) *Research on selected China's legal issues of e-business*. Springer, 23-32.
- Guthrie S. (2016) "These celebrities fell victim to a spam email". *The New Daily*, 16 March 2016. (accessed 9 August 2016, 8.30pm).
- Hackett R. (2016) "Everything you need to know about Google Allo's privacy backlash". *Fortune*, 22 September 2016. (accessed 13 February 2017, 10.03pm).
- Hafner K and Lyon M. (1996) *Where wizards stay up late: the origins of the Internet*, New York, US: Simon and Schuster.
- Harding L. (2016) "What are the Panama Papers? A guide to history's biggest data leak". *The Guardian*, 5 April 2016. (accessed 10 August 2016, 5.55pm).
- Hasinoff AA and Shepherd T. (2014) "Sexting in context: Privacy norms and expectations". *International Journal of Communication* 8: 2932-2955.
- Held V. (2005) "Feminist transformations of moral theory". In: Moore AD (ed) *Information Ethics: Privacy, Property and Power*. Seattle and London: University of Washington Press, 85-109.
- Helman L and Hannes S. (2016) "Tying executive compensation to privacy protection". *Unpublished Manuscript (in file with author)*.
- Henley J. (2005) "55 years on, the controversial kiss that could be worth £10,000". *The Guardian*, 13 April 2005. (accessed 30 March 2017, 9.44pm).

- Henry N, Powell A and Flynn A. (2017) "Not just 'revenge pornography': Australians' experience of image-based abuse". *RMIT University*, May 2017, rmit.edu.au. (accessed 8 May 2017, 3.58pm).
- Herman B. (1993) *The practice of moral judgment*, Cambridge, Mass., US: Harvard University Press.
- Hill K. (2011) "Fitbit moves quickly after users' sex stats exposed". *Forbes*, 5 July 2011. (accessed 12 August 2016, 6.07pm).
- Hobbs R and Jensen A. (2013) The past, present, and future of media literacy education. *Journal of media literacy education* 1, digitalcommons.uri.edu: 1-11.
- Hoffman C. (2010) "The battle for Facebook". *Rolling Stone*, 15 September 2010. (accessed 13 February 2017, 12.19pm).
- Holten E. (2015) "Consent, an objection". *Hysterical Feminisms*, 2015, hystericalfeminisms.com/consent. (accessed 17 July 2015, 5.39pm).
- Holten E, Jackson N, Bodker C, et al. (2015) "Someone stole naked pictures of me. This is what I did about it - video". *The Guardian*, 21 January 2015. (accessed 21 November 2016, 5.15pm).
- Hull G, Lipford HR and Latulipe C. (2011) "Contextual gaps: privacy issues on Facebook". *Ethics and Information Technology* 13: 289-302.
- iCloud. (2015) "iCloud terms and conditions". 16 September 2015. (accessed 16 August 2016, 6.31pm).
- Inness JC. (1992) *Privacy, intimacy, and isolation*, Oxford, UK: Oxford University Press.
- Jane E. (2016a) *Misogyny online: a short (and brutish) history*: SAGE Swifts.
- Jane EA. (2016b) "Online misogyny and feminist digilantism". *Continuum*: 1-14.
- Jenkins H. (2006) *Convergence culture: Where old and new media collide*, New York, US: NYU press.
- Jenkins Jr HW. (2010) "Google and the Search for the Future". *Wall Street Journal*, 14 August 2010. (accessed 26 July 2016 5.44pm).
- Johnson B. (2010) "Privacy no longer a social norm, says Facebook founder". *The Guardian*, 11 January 2010. (accessed July 26, 3.26pm).
- Johnson R. (2004) "Kant's moral philosophy". In: Zalta EN (ed) *The Stanford Encyclopedia of Philosophy*. plato.stanford.edu (Spring 2004 edition).

- Joye C. (2014) "Interview transcript: former head of the NSA and commander of the US cyber command, General Keith Alexander". *Financial Review*, 8 May 2014. (accessed 17 August 2016, 6.07pm).
- Juels A, Molnar D and Wagner D. (2005) "Security and privacy issues in e-passports". *First International Conference on Security and Privacy for Emerging Areas in Communications Networks (SECURECOMM'05)*. IEEE, 74-88.
- Kant I. (1870) *Grundlegung zur Metaphysik der Sitten*, (Ed. von Kirchmann, JH), Berlin, Germany: L. Heimann.
- Kant I. (1996a) *The metaphysics of morals*, (Trans. Gregor MJ), Cambridge, UK: Cambridge University Press.
- Kant I. (1996b) "On a supposed right to lie from philanthropy". In: Gregor MJ (ed) *Practical philosophy*. (Trans. Gregor MJ), Cambridge, UK: Cambridge University Press.
- Kant I. (1996c) "Toward perpetual peace". In: Gregor MJ (ed) *Practical philosophy*. (Trans. Gregor MJ), Cambridge, UK: Cambridge University Press.
- Kant I. (2009) *Groundwork of the Metaphysic of Morals*, (Trans. Paton HJ), New York, US: Harper Perennial Modern Thought.
- Kerstein S. (2009) Treating others merely as means. *Utilitas* 21: 163-180.
- Kiss J. (2014) "An online Magna Carta: Berners-Lee calls for bill of rights for web". *The Guardian*, 12 March 2014. (accessed 17 February 2017, 6.04pm).
- Kleingeld P. (2014) "The development of Kant's cosmopolitanism". In: Formosa P, Goldman A and Patrone T (eds) *Politics and teleology in Kant*. Cardiff, UK: University of Wales Press.
- Kleingeld P and Brown E. (2014) "Cosmopolitanism". In: Zalta EN (ed) *The Stanford Encyclopedia of Philosophy*. plato.stanford.edu (Fall 2014 edition).
- Komanduri S, Shay R, Norcie G, et al. (2011) "Adchoices? Compliance with online behavioral advertising notice and choice requirements". *ISJLP* 7: 603.
- Korsgaard CM. (1996) *Creating the kingdom of ends*, Cambridge, UK: Cambridge University Press.

- Kosinski M, Stillwell D and Graepel T. (2013) "Private traits and attributes are predictable from digital records of human behavior". *Proceedings of the National Academy of Sciences* 110: 5802-5805.
- Kramer AD, Guillory JE and Hancock JT. (2014) "Experimental evidence of massive-scale emotional contagion through social networks". *Proceedings of the National Academy of Sciences* 111: 8788-8790.
- Kranenborg H. (2015) "Google and the right to be forgotten". *European Data Protection Law Review* 1: 70-79.
- Kranzberg M. (1986) "Technology and History: 'Kranzberg's Laws' ". *Technology and culture* 27: 544-560.
- Kupfer J. (1987) "Privacy, autonomy, and self-concept". *American Philosophical Quarterly* 24: 81-89.
- Kurzweil R. (2005) *The singularity is near: When humans transcend biology*, New York, US: Penguin.
- Lanxing M. (2016) "The transparent self". *Ethics and Information Technology* 18: 9-16.
- Lardinois F. (2015) "Google launches native ads in Gmail to all advertisers". *Tech Crunch*, 1 September 2016. (accessed 11 August 2016, 6.06pm).
- Lawstuff. (2016) "Sexting". *Lawstuff*. (accessed 23 November 2016, 12.48pm).
- Lee M, Crofts T, McGovern A, et al. (2015) "Sexting and young people". *Report to the Criminology Research Advisory Council*, November 2015.
- LegislationUK. (2017) *Mental Capacity Act 2005*. legislation.gov.uk. (accessed 30 March 2017, 12.52pm).
- Leiner BM, Cerf VG, Clark DD, et al. (2009) "A brief history of the Internet". *ACM SIGCOMM Computer Communication Review* 39: 22-31.
- Lenhart A. (2009) Teens and sexting: how and why minor teens are sending sexually suggestive nude or nearly nude images via text messaging. *A Pew Internet & American Life Project Report*, 4. (accessed 10 August 2016, 2.29pm).
- Lessig L. (2006) *Code: Version 2.0*, New York, US: Basic Books.
- Levy A. (2016a) "Google parent Alphabet passes Apple market cap at the open". *CNBC*, 2 February 2016. (accessed 26 August 2016, 11.06pm).
- Levy S. (1984) *Hackers: Heroes of the computer revolution*, New York, US: Doubleday.

- Levy S. (2016b) "Why are we fighting the crypto wars again?". *Backchannel*, 11 March 2016. (accessed 19 August 2016, 5.50pm).
- Lewis K, Kaufman J and Christakis N. (2008) "The taste for privacy: an analysis of college student privacy settings in an online social network". *Journal of Computer - Mediated Communication* 14: 79-100.
- Lichtheim G. (1967) *The concept of ideology*, New York, US: Random House.
- LII. (2016) U.S. Constitution. *Legal Information Institute - Cornell University Law School*. (accessed 30 November 2016, 6.47pm).
- Lind J. (2004) Convergence: History of term usage and lessons for firm strategists. *Center for Information and Communications Research, Stockholm School of Economics, Working Paper*.
- Lindsay D. (2014) "The 'Right to Be Forgotten' in European Data Protection Law". In: Witzleb N, Lindsay D, Paterson M, et al. (eds) *Emerging challenges in privacy law: comparative perspectives*. Cambridge, UK: Cambridge University Press, pp. 290-337.
- Livingstone S and Hargrave AM. (2006) "Harmful to children? Drawing conclusions from empirical research on media effects". In: Carlsson U (ed) *Regulation, awareness, empowerment: young people and harmful media content in the digital age*. Göteborg, Sweden: The International Clearinghouse on Children, Youth and Media, 21-48.
- Lumby C. (2006) "Media ethics". In: Cunningham S and Turner G (eds) *The media and communications in Australia*. 2nd ed. Sydney, Australia: Allen & Unwin, 303-314.
- Lumby C and Albury K. (2010) "Too much? Too young? The sexualisation of children debate in Australia". *Media International Australia* 135: 141-152.
- Lupton D. (2014) "Self-tracking modes: Reflexive self-monitoring and data practices". *Imminent citizenships: personhood and identity politics in the informatic age*. ANU Canberra, 1-19.
- Lymn S. (2012) "Living in Orwell's world: how to disappear completely online". *The Conversation*, 3 April 2012. (accessed 22 August 2016, 4.43pm).
- Lynch J. (2015) "State courts strike blows to criminal DNA collection laws in 2014 - what to look for in 2015". *eff.org*, 5 January 2015. (accessed 23 November 2016, 4.50pm).

- Mackenzie C. (2014) "Three dimensions of autonomy: a relational analysis". In: Veltman A and Piper M (eds) *Autonomy, oppression and gender*. New York, US: Oxford University Press, 15-41.
- Mackenzie C and Stoljar N. (2000) *Relational autonomy: feminist perspectives on autonomy, agency, and the social self*, New York, US: Oxford University Press.
- MacKinnon R. (2012) "The netizen". *Development* 55: 201-204.
- Malheiros M, Jennett C, Patel S, et al. (2012) "Too close for comfort: a study of the effectiveness and acceptability of rich-media personalized advertising". *Proceedings of the SIGCHI conference on human factors in computing systems*. ACM, 579-588.
- Mann S. (2004) Sousveillance: inverse surveillance in multimedia imaging. *Proceedings of the 12th annual ACM international conference on Multimedia*. ACM, 620-627.
- Mansour RF. (2016) Understanding how big data leads to social networking vulnerability. *Computers in Human Behavior* 57: 348-351.
- Marwick AE. (2008) "To catch a predator? The MySpace moral panic". *First Monday*, 13.
- Matheson R. (2014) "A market for emotions: with emotion-tracking software, Affectiva attracts big-name clients, aims for 'mood-aware' Internet". *MIT News*, 31 July 2014. (accessed 21 September 2016, 2.01pm).
- Mathiesen T. (1997) "The viewer society: Michel Foucault's 'Panopticon' revisited". *Theoretical criminology* 1: 215-234.
- Matsui S. (2015) "The criminalization of revenge porn in Japan". *Washington International Law Journal* 24: 289.
- Mayer JR and Mitchell JC. (2012) "Third-party web tracking: Policy and technology". *2012 IEEE Symposium on Security and Privacy* IEEE, ieeexplore.ieee.org, 413-427.
- Mayer-Schönberger V. (2015) Keynote address: Privacy by regulation, 26 October 2015. *Amsterdam Privacy Conference*.
- Mayer-Schönberger V and Cukier K. (2013) *Big data: A revolution that will transform how we live, work, and think*, London, UK: Houghton Mifflin Harcourt.
- McGoogan C. (2016) "Twitter to verify more users: How to get a blue tick". *The Telegraph*, 20 July 2016. (accessed 12 August 2016, 3.46pm).

- McKendrick J. (2016) "Public cloud computing growing almost 50 per cent annually, Cisco says". *Forbes*, 31 May 2016, forbes.com. (accessed 1 March 2017, 6.40pm).
- McKenzie DF. (1984) "The sociology of a text: orality, literacy and print in early New Zealand". *The Library* 6: 333-365.
- Medeiros FA and Bygrave LA. (2015) "Brazil's Marco Civil da Internet: Does it live up to the hype?". *Computer Law & Security Review* 31: 120-130.
- Meese J. (2015) "Google Glass and Australian privacy law regulating the future of locative media". In: Wilken R and Goggin G (eds) *Locative Media*. New York, US: Routledge, 136-147.
- Meikle G and Young S. (2012) *Media convergence: networked digital media in everyday life*, London, UK: Palgrave Macmillan.
- Metz R. (2015) "Google Glass is dead; long live smart glasses". *MIT Technology Review* 118: 79-82.
- Michael K and Michael M. (2013) "The future prospects of embedded microchips in humans as unique identifiers: the risks versus the rewards". *Media Culture & Society* 35: 78-86.
- Michener J. (1999) System insecurity in the Internet age. *IEEE software* 16: 62-69.
- Mill JS. (2011) *On Liberty* (Project Gutenberg eBook). London, UK: The Walter Scott Publishing Co.
- Miller FH. (1987) "Informed consent for the man on the Clapham omnibus: an English cure for the American disease". *Western New England Law Review* 9 (1): 169-190.
- Miller N. (2015) "'Naked Rambler' Stephen Gough makes UK legal history by facing court in the nude". *smh.com.au*, 10 June 2015. (accessed 17 November 2016, 3.09pm).
- Miller N. (2016) "Behave, the Ministry of Nudge is watching you". *smh.com.au*, 15 September 2016. (accessed 15 September 2016, 10.45am).
- Millstein KH, Dare-Winters K and Sullivan S. (1994) "The power of silence: ethical dilemmas of informed consent in practice evaluation". *Clinical social work journal* 22: 317-329.
- Mislove A, Viswanath B, Gummadi KP, et al. (2010) "You are who you know: inferring user profiles in online social networks". *Proceedings of the third*

- ACM international conference on Web search and data mining*. ACM, 251-260.
- Mitchelson A. (2017a) "Peeping Tom drones prompt calls for a close look at privacy laws". *The New Daily*, 27 April 2017. (accessed 28 April 2017, 11.36am).
- Mitchelson A. (2017b) "WikiLeaks exposes the CIA's secret hacking tools". *The New Daily*, 8 March 2017, newdaily.com. (accessed 8 March 2017, 5.21pm).
- Mitrou L, Kandias M, Stavrou V, et al. (2014) "Social media profiling: A Panopticon or Omnipticon tool?". *Proc. of the 6th Conference of the Surveillance Studies Network*. 1-15.
- Mo JY. (2017) "Misuse of private information as a tort: The implications of *Google v Judith Vidal-Hall*". *Computer Law & Security Review* Vol. 33 (1) (Feb 2017): 87-97.
- Moor JH. (1990) "The ethics of privacy protection". *Library Trends* 39: 69-82.
- Moor JH. (1997) "Towards a theory of privacy in the information age". *Computers and Society* 27: 27-32.
- Moore AD. (2003) "Privacy: its meaning and value". *American Philosophical Quarterly* 40: 215-227.
- Moore AD. (2013a) "Privacy, speech, and the law". *Journal of Information Ethics* 22: 21-43.
- Moore AD. (2013b) "Privacy". In: LaFollette H (ed) *The International Encyclopedia of Ethics*. Wiley Online Library, 1-19.
- Moore S. (2004) "The doubling of place". In: Couldry N and McCarthy A (eds) *MediaSpace : place, scale, and culture in a media age*. London ; New York: Routledge, pp. 21-36.
- Morozov E. (2011) "Don't be evil". *The New Republic*, 13 July 2011. (accessed 24 August 2016, 11.01pm).
- Morris R. (2015) "Child watch: the apps that let parents 'spy' on their kids". *BBC News*, 29 January 2015. (accessed 10 August 2016, 4.24pm).
- Moser DJ, Schultz SK, Arndt S, et al. (2002) "Capacity to provide informed consent for participation in schizophrenia and HIV research". *American Journal of Psychiatry* 159: 1201-1207.

- Munro K. (2016) "Wenona bans social media app over concerns about risks to students". *smh.com.au*, 30 August 2016. (accessed 20 February 2017, 3.09pm).
- Murphy RF. (1964) "Social distance and the veil". *American Anthropologist* 66: 1257-1274.
- MyHealth. (2016) "Getting a My Health Record". *Australian Government: Australian Digital Health Agency*, 27 May 2016, myhealthrecord.gov.au. (accessed 24 April 2017, 10.14pm).
- Ness P. (2008) *The knife of never letting go*, London, UK: Walker Books.
- Nielsen. (2014) "Tech-styles: are consumers really interested in wearing tech on their sleeves?". *Nielsen Newswire*, 20 March 2014, nielsen.com. (accessed 24 February 2017, 11.01am).
- Niemietz. (1992) *Niemietz v. Germany*, Application No. 13710/88, Judgment of the European Court of Human Rights of 16 December 1992.
- Nightingale V. (2007) "New media worlds? Challenges for convergence". In: Nightingale V and Dwyer T (eds) *New media worlds: Challenges for convergence*. Oxford, UK: Oxford University Press, 19-36.
- Nimmer MB. (1954) "The right of publicity". *Law and Contemporary Problems* 19: 203-223.
- Nissenbaum H. (2010) *Privacy in context: Technology, policy, and the integrity of social life*, Stanford, California: Stanford University Press.
- Nissenbaum H. (2011) "A contextual approach to privacy online". *Daedalus* 140: 32-48.
- Nissenbaum H and Brunton F. (2015) *Obfuscation: A user's guide for privacy and protest*, Cambridge, Mass., US: MIT Press.
- Noack R. (2016) "This city embedded traffic lights in the sidewalks so that smartphone users don't have to look up". *Washington Post*, 25 April 2016. (accessed 31 May 2016 1.05pm).
- NSWSCLJ. (2016) *Remedies for the serious invasion of privacy in New South Wales*, 3 March 2016. New South Wales Standing Committee on Law and Justice.
- O'Neill O. (1989) *Constructions of reason: explorations of Kant's practical philosophy*, Cambridge, UK: Cambridge University Press.
- O'Neill O. (1996) *Towards justice and virtue: A constructive account of practical reasoning*, Cambridge, UK: Cambridge University Press.

- O'Neill O. (2002) *Autonomy and trust in bioethics*, Cambridge, UK: Cambridge University Press.
- O'Neill O. (2003) "Some limits of informed consent". *Journal of Medical Ethics* 29: 4-7.
- OAIC. (2017a) "Privacy Act", Office of the Australian Information Commissioner, oaic.gov.au. (accessed 3 February 2017, 4.25pm).
- OAIC. (2017b) "Privacy fact sheet 17: Australian Privacy Principles", Office of the Australian Information Commissioner, oaic.gov.au. (accessed 6 February 2017, 9.14pm).
- OAIC. (2017c) "Privacy law", Office of the Australian Information Commissioner, oaic.gov.au. (accessed 2 February 2017, 1.04pm).
- Obar JA and Oeldorf-Hirsch A. (2016) "The biggest lie on the internet: ignoring the privacy policies and terms of service policies of social networking services". *TPRC 44: The 44th Research Conference on Communication, Information and Internet Policy 2016*.
- OECD. (1980) "OECD Privacy Principles". oecdprivacy.org. (accessed 16 February 2017, 1.42pm).
- Ofcom. (2016) Adults' media use and attitudes report 2016.
- OHCHR. (1976) International Covenant on Civil and Political Rights. United Nations Human Rights Office of the High Commissioner, ohchr.org. (accessed 15 February 2017, 7.24pm).
- OHCHR. (2017) "Status of ratification: interactive dashboard". United Nations Human Rights Office of the High Commissioner, indicators.ohchr.org. (accessed 15 February 2017, 7.30pm).
- OJEU. (2016) "Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)". *Official Journal of the European Union*, eur-lex.europa.eu. (accessed 3 February 2017, 1.12pm).
- Olding R. (2014) "Spyware's role in domestic violence". *smh.com.au*, 22 March 2014. (accessed 10 August 2016, 4.29pm).
- Olding R and Munro P. (2016) "The day Laura Pilati saw her photo on a schoolgirl porn site". *smh.com.au*, 19 August 2016. (accessed 22 August 2016, 5.09pm).

- Orwell G. (1949) *1984*, London, UK: Penguin.
- Parent WA. (1983) "Privacy, morality, and the law". *Philosophy & Public Affairs*: 269-288.
- Parfit D. (2011) *On what matters: volume one*, Oxford, UK: Oxford University Press.
- Pariser E. (2011) *The filter bubble: What the Internet is hiding from you*: Penguin UK.
- Pateman C. (1980) "Women and consent". *Political Theory* 8: 149-168.
- Patterson H. (2013) "Contextual expectations of privacy in self-generated health information flows". *TPRC*, 1-48.
- Perera C, Ranjan R, Wang L, et al. (2015) Big data privacy in the internet of things era. *IT Professional* 17: 32-39.
- Perez E, Prokupecz S and Cohen T. (2014) "More than 90 people nabbed in global hacker crackdown". *CNN*, 20 May 2014. (accessed 9 August 2016, 6.03pm).
- Perlroth N. (2016) "Apple rushes out iPhone software fix to avoid conversation eavesdropping". *smh.com.au*, 26 August 2016. (accessed 26 August 2016, 5.43pm).
- Perlroth N. (2017) "Spyware's odd targets: backers of Mexico's soda tax". *The New York Times*, 11 February 2017. (accessed 14 February 2017, 3.21pm).
- Pesce JP, Casas DL, Rauber G, et al. (2012) "Privacy attacks in social media using photo tagging networks: a case study with Facebook". *Proceedings of the 1st Workshop on Privacy and Security in Online Social Media*. ACM, 1-8.
- Peters J. (2016) *The Idealist: Aaron Swartz and the Rise of Free Culture on the Internet*, New York, US: Simon and Schuster.
- Peterson A. (2016) "Why a staggering number of Americans have stopped using the Internet the way they used to". *The Washington Post*, 13 May 2016. (accessed 19 August 2016, 7.01pm).
- Pilgrim T. (2015) Face-to-face interview with Sacha Molitorisz, 8 October 2015.
- Piwek L and Joinson A. (2016) "'What do they Snapchat about?' Patterns of use in time-limited instant messaging service". *Computers in Human Behavior* 54: 358-367.
- Poikola A, Kuikkaniemi K and Honko H. (2010) "MyData: a Nordic model for human-centered personal data management and processing". Ministry of

- Transport and Communication, Finland, lvm.fi. (accessed 14 February 2017, 1.52pm).
- Posner RA. (1977) "The right of privacy". *Georgia Law Review* 12.
- Price R. (2014) "These videos demonstrate exactly what is wrong with Google Glass". *The Daily Dot*, 28 March 2014. (accessed 10 August 2016, 5.50pm).
- Prosser WL. (1960) "Privacy". *California Law Review* 48: 383-423.
- Putzer GJ and Park Y. (2010) "Effects of innovation factors on smartphone adoption by nurses in community hospitals". *Perspectives in Health Information Management*.
- quantifiedself.com. (2016) "Quantified Self". quantifiedself.com. (accessed 12 August 2016, 6.04pm).
- Rachels J. (1975) "Why privacy is important". *Philosophy & Public Affairs* 4: 323-333.
- Rauhofer J. (2008) "Privacy is dead, get over it! Information privacy and the dream of a risk-free society". *Information & Communications Technology Law* 17: 185-197.
- Rauscher F. (2016) "Kant's social and political philosophy". In: Zalta EN (ed) *The Stanford Encyclopedia of Philosophy*. plato.stanford.edu (Fall 2016 edition).
- Reiman J. (1976) "Privacy, intimacy, and personhood". *Philosophy & Public Affairs* 6: 26-44.
- Reiman J. (2004) "Driving to the panopticon: A philosophical exploration of the risks to privacy posed by the information technology of the future". In: Rössler B (ed) *Privacies: Philosophical Evaluations*. Stanford, US: Stanford University Press, 194-214.
- Richards NM and Solove DJ. (2010) "Prosser's Privacy Law: A Mixed Legacy". *California Law Review* 98 (6): 1887-1924.
- Rigney S. (2017) "Murder trial almost aborted after witness sends Facebook friend request to juror". *smh.com.au*, 9 February 2017. (accessed 10 February 2017, 5.31pm).
- Roberts D and Ackerman S. (2015) "NSA mass phone surveillance revealed by Edward Snowden illegal". *The Guardian*, 8 May 2015. (accessed 6 February 2017, 5.26pm).
- Ronson J. (2015) *So you've been publicly shamed*, London, UK: Picador.

- Rosen J. (2003) "PressThink: An introduction". *PressThink*. (accessed 3 August 2016, 3.29pm).
- Rosen J. (2006) "The people formerly known as the audience". *PressThink*. (accessed 3 August 2016, 3.35pm).
- Rosenzweig R. (2004) "How will the net's history be written? Historians and the Internet". In: Nissenbaum H and Price ME (eds) *Academy and the Internet*. New York, US: Peter Lang.
- Rössler B. (2004) *Privacies: philosophical evaluations*, Stanford, US: Stanford University Press.
- Rössler B. (2005) *The value of privacy*, Malden, Mass, US: Polity.
- Rotenberg M and Jacobs D. (2013) "Updating the law of information privacy: the new framework of the European union". *Harvard Journal of Law and Public Policy* 36: 605-652.
- Rubinfeld J. (1989) "The right of privacy". *Harvard law review* 102: 737-807.
- Rudesill DS, Caverlee J and Sui D. (2015) "The deep web and the darknet: A look inside the internet's massive black box". *Woodrow Wilson International Center for Scholars, STIP*, 3. (accessed 2 August 2016, 1.51pm).
- Rush E. (2012) "Children, media and ethics". In: Warburton W and Braunstein D (eds) *Growing up fast and furious: Reviewing the impacts of violent and sexualised media on children*. Sydney, Australia: Federation Press, 159-174.
- Rushkoff D and Goodman B. (2004) "The Persuaders". *Frontline*. (accessed 9 August 2016, 2.54pm).
- Sandel MJ. (2012) *What money can't buy: the moral limits of markets*, London, UK: Penguin.
- Sarigol E, Garcia D and Schweitzer F. (2014) "Online privacy as a collective phenomenon". *Proceedings of the second ACM conference on Online social networks*. ACM, 95-106.
- Sartre J-P. (2005) *Huis clos*, Oxford, UK: Routledge.
- Scannell P. (1996) "Dailiness". *Radio, television and modern life : a phenomenological approach*. Oxford, UK; Cambridge, Mass., USA: Blackwell, pp. 144-178.
- Schmidt E and Cohen J. (2013) *The new digital age: Reshaping the future of people, nations and business*: Hachette UK.

- Schneier B. (2013) "NSA surveillance: a guide to staying secure." *The Guardian*, 6 September 2013. (accessed 22 August 2016, 3.08pm).
- Schneier B. (2015) *Data and Goliath: The hidden battles to collect your data and control your world*, New York, US; London, UK: WW Norton & Company.
- Schoeman F. (1984) "Privacy: philosophical dimensions". *American Philosophical Quarterly* 21: 199-213.
- Schrems. (2015) *Maximillian Schrems v Data Protection Commissioner*, C-362/14, Judgment of the Court of Justice of the European Union (Grand Chamber) of 6 October 2015. (accessed 1 December 2016, 12.54pm).
- Scott E. (2015) "Senate passes controversial metadata laws". *The Sydney Morning Herald*, 27 March 2015. (accessed 27 July 2016, 5.01pm).
- Scott GG. (2014) "More than friends: popularity on Facebook and its role in impression formation". *Journal of Computer - Mediated Communication* 19: 358-372.
- Segall L. (2015) "Pastor outed on Ashley Madison commits suicide". *CNN.com*, 8 September 2015. (accessed 3 February 2017, 4.57pm).
- Sengupta S. (2013) "What you didn't post, Facebook may still know". *The New York Times*, 26 March 2013. (accessed 9 August 2016, 3.56pm).
- Shankar S, Udupi V and Gavas RD. (2016) "Biometric verification, security concerns and related issues - a comprehensive study". *International Journal of Information Technology and Computer Science* 2016 (4): 42-51.
- Sherman N. (1997) *Making a necessity of virtue: Aristotle and Kant on virtue*, Cambridge, UK: Cambridge University Press.
- Shih G. (2013) "Boston Marathon bombings: how Twitter and Reddit got it wrong". *The Independent*, 21 April 2013. (accessed 26 August 2016, 3.46pm).
- Shvartzshnaider Y, Tong S, Wies T, et al. (2016) "Learning privacy expectations by crowdsourcing contextual informational norms". *cs.nyu.edu*. (accessed 21 February 2017, 2.30pm).
- Simonite T. (2015) "How Facebook and Google's plans to boost internet access advanced in 2015". *MIT Technology Review*, 24 December 2015. (accessed 2 August 2016, 12.34pm).

- Singer N. (2012) "Acxiom, the quiet giant of consumer database marketing". *The New York Times*, 16 June 2012. (accessed 9 August 2016, 2.59pm).
- Skegg P. (2011) "Presuming competence to consent: could anything be sillier". *University of Queensland Law Journal* 30 (2): 165-187.
- Sklar J. (2015) "Hired and fired by algorithm". *MIT Technology Review* 118: 6: 69.
- Sleeper M, Cranshaw J, Kelley PG, et al. (2013) "I read my Twitter the next morning and was astonished: a conversational perspective on Twitter regrets". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 3277-3286.
- Slefo G. (2016) "Pinterest partners with Oracle to measure offline sales". *Advertising Age*, 7 June 2016. (accessed 15 August 2016, 5.27pm).
- Smith C. (2010) "7,500 online shoppers accidentally sold their souls to Gamestation". *The Huffington Post*, 17 June 2010. (accessed 5 April 2016, 6.27pm).
- Smith M. (1987) "The Humean theory of motivation". *Mind* 96: 36-61.
- Smith M. (1994) *The moral problem*, Oxford, UK: Wiley-Blackwell.
- Smith MR and Marx L. (1994) *Does technology drive history? The dilemma of technological determinism*, Cambridge, Mass., UK: MIT Press.
- Snow HA and Fleming BR. (2014) "Consent, capacity and the right to say no". *The Medical journal of Australia* 201: 486-488.
- Sobel DL. (2000) "The process that 'John Doe' is due: addressing the legal challenge to internet anonymity". *Virginia Journal of Law and Technology*, 5, vjolt.net. (accessed 19 August 2016, 6.14pm).
- SocialYup. (2016) "Buy Facebook friends". *SocialYup.com*. (accessed 30 August 2016, 8.31pm).
- Solove DJ. (2007) "'I've got nothing to hide' and other misunderstandings of privacy". *San Diego law review* 44: 745-772.
- Solove DJ. (2008) *Understanding privacy*, Cambridge, Mass., US: Harvard University Press.
- Spinello RA. (2011) "Privacy and social networking technology". *International Review of Information Ethics* 16: 41-46.
- Stallman RM. (2002) "Chapter 4: why software should not have owners". *Free software, free society: selected essays of Richard M. Stallman*. Boston, US: The Free Software Foundation, 47-51.

- Standen PJ, Brown DJ, Battersby S, et al. (2011) "A study to evaluate a low cost virtual reality system for home based rehabilitation of the upper limb following stroke". *International Journal on Disability and Human Development* 10: 337-341.
- Stanley B, Guido J, Stanley M, et al. (1984) "The elderly patient and informed consent: empirical findings". *Jama* 252: 1302-1306.
- Stasko E and Geller P. (2015) "Reframing sexting as a positive relationship behavior". *American Psychological Association*. (accessed 10 August 2016, 2.41pm).
- Statistica. (2016) "Number of monthly active Facebook users worldwide as of 2nd quarter 2016". *Statistica*. (accessed 12 August 2016, 3.08pm).
- Stoycheff E. (2016) "Under surveillance: examining Facebook's spiral of silence effects in the wake of NSA internet monitoring". *Journalism & Mass Communication Quarterly* 93: 296-311.
- Strassberg DS, McKinnon RK, Sustaita MA, et al. (2013) Sexting by high school students: An exploratory and descriptive study. *Archives of Sexual Behavior* 42: 15-21.
- Streeter T. (1999) "That deep romantic chasm': Libertarianism, neoliberalism, and the computer culture". In: Calabrese A and Burgelman J-C (eds) *Communication, citizenship, and social policy: rethinking the limits of the welfare state*. . Lanham, Maryland, US: Rowman & Littlefield, 49-64.
- Stuart K. (2014) "Brianna Wu and the human cost of Gamergate: 'Every woman I know in the industry is scared'". *The Guardian*, 18 October 2014. (accessed 10 August, 5.00pm).
- Summers CA, Smith RW and Reczek RW. (2016) "An audience of one: Behaviorally targeted ads as implied social labels". *Journal of Consumer Research*. (accessed 16 August 2016, 12.15pm).
- Sweeney L. (2000) "Simple demographics often identify people uniquely". *Carnegie Mellon University, Data Privacy Working Paper 3*. (accessed 19 August 2016, 12.56pm).
- Taddei S and Contena B. (2013) "Privacy, trust and control: Which relationships with online self-disclosure?". *Computers in Human Behavior* 29: 821-826.
- Taddicken M. (2014) "The 'privacy paradox' in the social web: the impact of privacy concerns, individual characteristics, and the perceived social

- relevance on different forms of self-disclosure". *Journal of Computer-Mediated Communication* 19: 248-273.
- Tashea J. (2017) "Courts are using AI to sentence criminals. That must stop now". *Wired*, 17 April 2017, wired.com. (accessed 18 April 2017, 11.31am).
- Tavani HT. (2007) "Philosophical theories of privacy: implications for an adequate online privacy policy". *Metaphilosophy* 38: 1-22.
- Tene O and Polonetsky J. (2015) "A theory of creepy: technology, privacy, and shifting social norms". *Yale Journal of Law and Technology* 16: 58-102.
- Thomas J. (2015) "Hurtcore porn and the dark web: why we need an ethics of technology". *BigThink*, 19 September 2015. (accessed 26 August 2016, 4.11pm).
- Thomson JJ. (1975) "The right to privacy". *Philosophy & Public Affairs*: 295-314.
- Titcomb J. (2016) "Why has Mark Zuckerberg taped over the webcam and microphone on his MacBook?". *The Telegraph*, 23 June 2016. (accessed 19 August 2016, 7.04pm).
- TNSOED. (1993) In: Brown L (ed) *New Shorter Oxford English Dictionary on Historical Principles*. Oxford, UK: Clarendon.
- Tomlinson J. (2007) *The culture of speed: the coming of immediacy*, London, UK: Sage.
- TrackMeNot. (2016), cs.nyu.edu/trackmenot. (accessed 22 August 2016, 4.30pm).
- Trottier D. (2016) "Digital vigilantism as weaponisation of visibility". *Philosophy & Technology*: 1-18.
- Tsukayama H. (2012) "FTC announces \$22.5 settlement with Google". *The Washington Post*, 9 August 2012, washingtonpost.com. (accessed 10 April 2017, 3.41pm).
- Turkle S. (2011) *Alone together: Why we expect more from technology and less from each other*, New York, US: Basic books.
- Turner A. (2016) "Google Translate becomes more polyglot". *smh.com.au*, 1 June 2016. (accessed 17 August 2016, 2.55pm).
- Turow J, Hennessy M and Draper NA. (2015) "The tradeoff fallacy: how marketers are misrepresenting American consumers and opening them up to exploitation". *A report from the Annenberg School for Communication*. (accessed 28 October 2016, 6.12pm).

- Turow J, Hoofnagle CJ, Mulligan DK, et al. (2007) "The Federal Trade Commission and consumer privacy in the coming decade". *ISJLP* 3: 723-749.
- Twitter. (2017) "Twitter Privacy Policy: information collection and use - third-parties and affiliates". twitter.com. (accessed 5 April 2017, 9.47pm).
- UN. (1948) United Nations Universal Declaration of Human Rights. (accessed 1 December 2016, 2.41pm).
- UN. (2015) "The state of broadband 2015: broadband as a foundation for sustainable development". September 2015, broadbandcommission.org. (accessed 1 December 2016, 2.43pm).
- UN. (2016) "The state of broadband: broadband catalyzing sustainable development, September 2016", United Nations' Broadband Commission for Digital Development. September 2016, broadbandcommission.org. (accessed 2 March 2017, 3.15pm).
- UNGA. (2016) "The promotion, protection and enjoyment of human rights on the Internet", Human Rights Council, United Nations General Assembly, article19.org. 27 June 2016. (accessed 7 February 2017, 1.14pm).
- Vallor S. (2012) "Social networking and ethics". In: Zalta EN (ed) *The Stanford Encyclopedia of Philosophy*. plato.stanford.edu (Winter 2016 edition).
- Vallor S. (2016) *Technology and the virtues: a philosophical guide to a future worth wanting*, New York, US: Oxford University Press.
- van Dijk J. (2006) *The network society*, London, UK: Sage Publications.
- Veltman A and Piper M. (2014) "Introduction". In: Veltman A and Piper M (eds) *Autonomy, oppression and gender*. New York, US: Oxford University Press, 1-11.
- Victoria. (2016) "Failure to disclose offence". *Victoria State Government Justice and Regulation*. (accessed 18 November 2016, 6.13pm).
- Virilio P. (2002) "The visual crash". In: Levin TY, Frohne U and Weibel P (eds) *CTRL [Space]*. 2nd ed. Cambridge, Mass, US: MIT Press, 108-113.
- Wang D, Park S and Fesenmaier DR. (2012) "The role of smartphones in mediating the touristic experience". *Journal of Travel Research* 51: 371-387.
- Wang Y, Leon PG, Acquisti A, et al. (2014) "A field trial of privacy nudges for Facebook". *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*. ACM, 2367-2376.

- Warren SD and Brandeis LD. (1890) "The right to privacy". *Harvard law review* Vol. 4, No. 5: 193-220.
- Waters N. (2015) Phone interview with Sacha Molitorisz, 10 December 2015.
- Welch M. (2013) *Corrections: A critical approach*, London, UK: Routledge.
- Wen P. (2016) "China launches 'hack-proof' quantum satellite in world first". *smh.com.au*, 17 August 2016. (accessed 22 August 2016, 3.36pm).
- Westin AF. (1967) *Privacy and freedom*, New York, US: Atheneum.
- Whitman JQ. (2004) "The two western cultures of privacy: dignity versus liberty". *Yale Law Journal* Vol. 113, No. 6: 1151-1221.
- Wikipedia. (2016a) "Internet". *Wikipedia*. (accessed 18 April 2016, 12.22pm).
- Wikipedia. (2016b) "Welcome to Wikipedia". *Wikipedia*. (accessed 3 August 2016, 5.40pm).
- Williams G. (2016) "Kant's account of reason". In: Zalta EN (ed) *The Stanford Encyclopedia of Philosophy*. plato.stanford.edu (Spring 2016 edition).
- Wood A. (2014) "Kant's principles of publicity". In: Formosa P, Goldman A and Patrone T (eds) *Politics and teleology in Kant*. Cardiff, UK: University of Wales Press.
- Wood AW. (1999) *Kant's ethical thought*, Cambridge, UK: Cambridge University Press.
- Woodlock D. (2016) "The abuse of technology in domestic violence and stalking". *Violence against women*, 12 May 2016. (accessed 10 August 2016, 4.39pm).
- WorldWideWebSize. (2016) "World Wide Web Size". (accessed 2 August 2016, 1.19pm).
- Wright D and Raab C. (2014) "Privacy principles, risks and harms". *International Review of Law, Computers & Technology* 28: 277-298.
- Yadron D. (2016) "Facebook planning encrypted version of its Messenger bot, sources say". *The Guardian*, 31 May 2016. (accessed 22 August 2016, 2.48pm).
- Yorke S. (2015) Email interview with Sacha Molitorisz, 3 November 2015.
- Young E. (2015) "Educational privacy in the online classroom: FERPA, MOOCs, and the big data conundrum". *Harvard Journal of Law and Technology* Vol. 28, No. 2, Spring 2015: 549-592.

- Young IM. (2004) "A room of one's own: old age, extended care, and privacy". In: Rössler B (ed) *Privacies: Philosophical Evaluations*. Stanford, US: Stanford University Press, 168-186.
- Young-Bruehl E. (2008) *Why Arendt matters*, New Haven, Conn., US: Yale University Press.
- Yuhas A. (2016) "Hacker who stole nude photos of celebrities gets 18 months in prison". *The Guardian*, 28 October 2016. (accessed 22 February 2017, 4.24pm).
- Zang J, Dummit K, Graves J, et al. (2015) "Who knows what about me? A survey of behind the scenes personal data sharing to third parties by mobile apps". *Technology Science* 30 October 2015: 1-53.
- Zara C. (2015) "Facebook privacy lawsuit alleges social network created 'shadow profiles' of non-users". *International Business Times*, 24 April 2015. (accessed 11 August 2016, 6.52pm).
- Zetter K. (2016) "Hacker lexicon: what are white hat, gray hat, and black hat hackers?". *Wired*, 13 April 2016. (accessed 24 August 2016, 5.54pm).
- Zimmer M and Hoffman A. (2012) "Privacy, context, and oversharing: reputational challenges in a web 2.0 world". In: Masum H, Tovey M and Newmark C (eds) *The reputation society: How online opinions are reshaping the offline world*. London, UK; Cambridge, Mass., US: The MIT Press, 175-184.
- Zittrain J. (2008) *The future of the internet - and how to stop it*, New Haven, Conn., US: Yale University Press.