1

On the Project and Its Motivation

1.1 THE PROJECT: USHERING IN THE AGE OF SURVEILLANCE

We are all targets in the age of surveillance.[1] Imagine that you buy the following items: Cocoa-butter lotion, a large purse, vitamin supplements (zinc and magnesium) and a bright blue rug. Now imagine that these purchases tell someone that there is an 87 per cent chance that you are pregnant and due in six months. The company Target did this – used these and other pieces of data to produce a 'pregnancy score' for its customers. This surveillance became public knowledge when Target sent a family a 'pregnancy pack', congratulating the shopper on their pregnancy. The shopper was a teenage girl, living at home and her parents were unaware that their child was sexually active, let alone pregnant (Hill, 2012). Having trivial information such as purchase of cocoa butter produce Personal Information such as pregnancy is a fact of the age of surveillance. Coupling our behaviour with this focused attention reveals information for those who would want to target us. And while many would see such use of Personal Information as problematic, explaining why the purchase of cocoa butter is something of deep moral importance is far harder and more complex.

Information communication technologies (ICTs) have revolutionised the ways we live our lives. They are ubiquitous – firmly integrated into our working habits and our social lives – and play an ever deeper role in the exercise of basic political rights. From its initial introduction to the public in the early-mid 1990s, internet access has become

[1] I will not seek to define surveillance, as opening with a definitional discussion can tend to obscure the larger points being made. That said, this description by David Lyon will serve as a functionally useful account of how I generally use the term. 'Literally, surveillance means to "watch over" and as such it is an everyday practice in which human beings engage routinely, often unthinkingly … In most instances, however, surveillance has a more specific usage, referring to some focused and purposive attention to objects, data, or persons' (Lyon, 2009). My interest is with this sustained and focused attention to a target or set of targets. This is returned to in Chapter 3, where I talk about surveillance operators as epistemic actors.

comprehensive in many developed countries (The World Bank 2015)[2] to such a point that the UN Special Rapporteur on the Promotion and Protection of the Right to Freedom of Opinion and Expression held that internet access could be considered a basic human right (The Human Rights Council 2011). Social networking defines many people's social lives and is used by researchers, marketers and governments to shed light on and assist in, provision of essential services like healthcare (Chambers et al., 2012; Househ, 2013; Moorhead et al., 2013).

Our behaviours are evolving in parallel with these ubiquitous technologies. Whether it is the drive to photograph and broadcast our meals, the role played by Twitter in the social uprisings in the Middle East in 2011 (Lotan et al., 2011),[3] the push to have police officers wear body cameras while on patrol (Belluck, 2009; Pearl, 2015) or the live webcasting of brain surgery (Belluck, 2009), we use ICTs to communicate a wealth of Personal Information. Underneath the ubiquity of ICTs is the huge range of different information technology types, connected through their capacity to produce, collect, store and communicate information. As David Lyon notes, the evolution of our social practices runs hand in hand with the development of surveillance technologies:

> Although as a set of practices it is as old as history itself, systematic surveillance became a routine and inescapable part of everyday life in modern times and is now, more often than not, dependent on information and communication technologies (ICTs). Indeed, it now makes some sense to talk of 'surveillance societies', so pervasive is organizational monitoring of many kinds.
>
> (Lyon, 2009)

Combine the near invisible presence of ICTs in our lives with their informational capacities and we have *the age of surveillance*: a social epoch marked by informational technologies which endorse, encourage and enable us to live lives under constant surveillance.

What marks this age as one of *surveillance* is our own role in this – it is not simply that there are these new information technologies that target us for observation. We are complicit in this observation – we are often the willing sources of this information, happily uploading selfies, buying wearable surveillance technologies, actively publicising vast amounts of Personal Information like no other time in history. These ICTs are not just invasive; they are changing our very behaviours. What's so unique is that ICTs afford the ability to make ourselves the subject of observation. Facebook's value comes from the fact that its users are the active suppliers of Personal Information. This involvement in our own surveillance is unique in history.

[2] Note here that such coverage and access is patchy both globally and within different sectors of the community.

[3] At least in its initial stages. As Evgeny Morozov notes, social media alone was not enough to continue and close off a full revolution (Morozov, 2013, pp. 127–128).

What is interesting about this age of surveillance is our ambivalence to the treatment of this Personal Information. Compare the responses to a new smart phone with the revelations of widespread government surveillance: when Apple released their iPhone6, a colleague told me glowingly that it will remind him of things on his shopping list, can use the GPS to send text messages to his wife when he's almost home from work and can monitor his patterns of sleep. In the same conversation, he spoke with anger of how various governments around the world have been exposed for widespread spying and ubiquitous surveillance (Greenwald, 2014; Harding, 2014). His concern about government overreach seemed justified, but it was hard to reconcile this with his willing encouragement of commercial technology that watches him while he sleeps. Imagine if the government was involved in surveillance at this level of intimacy, watching him while he sleeps. This ambivalence is confusing – should we continue to embrace these surveillance technologies, further involving them in our most intimate of behaviours? Or should we be deeply worried, offended at the very thought of strangers watching us at every turn?

At first glance, our ambivalent responses suggest revealed preferences:[4] though we say we are opposed to such widespread surveillance, our behaviours reveal that we don't actually care so much. If anything, our willing *involvement* in surveillance shows that we actively endorse this age. Part of the explanation for this is that so much of this Personal Information is *insignificant*: who honestly cares whether I roll over in bed at 3:48am on Monday the 10th of August? Any worry about such insignificant information is likely to be paranoid and self-obsessed. However, as the responses to Edward Snowden's revelations of government surveillance show (Greenwald, 2014), there are deep concerns for many people about such omnipresent surveillance.

A sustained ethical analysis of these ICTs and of our behaviours surrounding them, will show that the arguments of involvement and insignificance are wrong-footed: by looking more deeply at just what Personal Information is and the ways that ICTs produce, collect, store and communicate this Personal Information, we can recognise the rise of *Virtual Identities*. These Virtual Identities carry with them special moral importance. As we will see, framing the discussion of surveillance in relation to Virtual Identities not only explains *why* our concerns are justified, it also gives some direction as to *how* to respond to those concerns, in ways that allow us to retain many of the desirable aspects of these ICTs. By looking at the age of surveillance in reference to Virtual Identities, we can both better understand the limits of the arguments of involvement and insignificance and the ambivalence to these technologies.

[4] As Daniel Solove points out, the arguments that our actions betray our real preferences are quite weak when looked at closely (Solove, 2004, pp. 76–92).

1.2 THE MOTIVATION: FROM INTIMATE TO INNOCUOUS INFORMATION

The motivation of this book is to understand how we should be treating Personal Information. Our ambivalence to Personal Information presents a challenge: should we care about privacy and control over Personal Information or perhaps there simply is no real moral problem – the sorts of information that surveillance technologies produce are innocuous, of no moral weight. Any worry about such information is mistaken. However, if we think of intimate Personal Information such as medical information, we see that there are legitimate concerns about the collection, use, storage and communication of medical information. If Personal Information collected by surveillance technologies shares some trait with medical information, then perhaps we ought to be taking more care with Personal Information more generally.

The differential uptake between government and private information services illustrates a general public ambivalence to surveillance technologies. In the past decade, many countries have attempted to roll out electronic patient medical records, but have had limited success. Australia passed the *Healthcare Identifiers Act 2010* (Australian Government 2010), described by the then Minister for Health Care and Ageing, Nicola Roxon, as 'a key building block of the Government's plans ... to revolutionise healthcare delivery through the introduction of personally-controlled electronic health records' (Nicola Roxon, 2010). The Australian eHealth programme began in July 2012.[5] To date, this 'key building block' has had limited uptake by the community (Ley, 2015). In early 2009, US president Barack Obama committed to 'a goal of computerizing all of America's medical records within 5 years as a means of improving efficiency, quality and safety and ultimately saving money' (Tang and Lee, 2009). The roll out in the United States has been plagued with a host of problems (Bowman, 2013). While some efforts have been more effective than others, the overall trend seems to be limited community engagement in these government initiatives.

At the same time, private industry is releasing an increasing range of services and products that collect, store, analyse and distribute Personal Information. Wearable health informatics technologies like Fitbit promise to 'empower and inspire you to live a healthier, more active life. We design products and experiences that fit seamlessly into your life so you can achieve your health and fitness goals, whatever

[5] As of July 2012, the Australian eHealth programme has been released, see ehealth.gov.au. However, it is operating on an opt-in system to enrol users and as of the date of writing, has been unsuccessful in getting people to enrol in the programme. While this lack of enrolment is due to a host of reasons, it certainly indicates considerable scepticism in such projects – including whether the Australian people trust the government with their data and whether such programmes will actually be useful or not. It should be noted that these eHealth programmes have generally been problematic: the United Kingdom attempt to convert medical records to electronic formats, proposed in 2002, had cost £10 billion by 2013 and was largely abandoned (Syal, 2013).

they may be' (FitBit, 2015). They produce information about the wearer's physical activity, in order to better inform the individual wearer about their lifestyle. As the technologies have advanced, 'informational ecosystems' have evolved into a suite of technologies and services. Apple's informational services are being bolstered by physical instruments such as the iPhone, iPad and the iWatch. The iRing (currently under patent application) involves 'an advanced ring-style wearable that uses voice, motion and touch input to control and interact with larger computing devices' (Campbell, 2015). What's notable about information harvesting products and services like those released and offered by Apple is that people actively seek them out and in many cases will queue up for hours or days to be the first to purchase the product. This is no minor subsection of the community either – Apple's position as the world's most highly valued company is built on the success of their informational ecosystems. Compared to government-provided services, these personalised information devices show people's willingness and genuine excitement to engage in self surveillance.

Perhaps recognising this ambivalence, the United Kingdom has sought to open up public access to medical data in order to stimulate new ways of producing and using aggregated medical data and new ways of producing data that is relevant for healthcare. Vast amounts of anonymised National Health Service (NHS) data have been made available for public access as of September 2012.[6]

One response is to point out the difference between government and private actors by saying that government services are about intimate *medical* information, whereas commercial products and services simply use innocuous Personal Information – we're talking about apples and oranges here. However, whether governmental, private or a mixture[7] of both (Solove, 2004, p. 3), a key goal of the informationalisation of our lives is to *integrate* different information to produce new information.[8] In addition to making medical data easier to communicate and access, what's essential to recognise is that by integrating this innocuous information, surveillance technologies are expanding what qualifies as morally relevant information.

To show how innocuous information becomes morally weighty, consider the problem of aging populations in the developed world. As the number and proportion of aged people in populations increase, dementia is also likely to increase in

[6] A summary of the motivations outlined in §1.1 of a recent UK Government White Paper, where it is stated that the release of government data is 'enabling people to make better choices about the public services they use and to hold government to account on spending and outcomes. Transparency is also providing the raw material for innovative new business ventures and for public service professionals to improve their performance.' (U.K. Government 2012).

[7] I refer here to the idea that a clear distinction between international state actors and global non-state actors is not a clear distinction. This point is borne out by the third motivation of the United Kingdom's Open Health programme, which is but one example of an expressly public/private partnership involving the use of personal information.

[8] The notion of new information and emergence is discussed in detail in Chapter 5.

number of sufferers and possibly in severity (Malloy et al., 2007, pp. 77–78). Remote patient monitoring and early detection may not only decrease economic costs (Tegart, 2010, pp. 8–9) but can also, hopefully, increase the quality of life for sufferers and carers:

> Near the end, my parents were spending about $180 day for home nursing. For just a fraction of their monthly nursing bill, they could have thrown enough blinking sensors and networking gizmos into their house to record and transmit every step, bite, breath, word and heartbeat in their Portland house.
>
> (Baker, 2007, p. 157)

As the need for treatment and support of those with conditions like dementia increases, there will likely be a corresponding growth in markets interested in exploiting these opportunities by developing novel ways and means of identifying and treating sufferers and supporting their carers.[9]

Personal Information, understood as 'information that relates to a person or group of people in some way'[10] that is non-intimate and innocuous, is a key tool in strategies to mitigate the impacts of aging populations. Consider the length of time between hearing a close friend or relative's voice on a phone and the recognition of who is speaking. This lapse between hearing the voice and recognising who is speaking is being investigated as a potential flag for dementia in elderly people (Baker, 2007, p. 168). Another novel method looks at word and grammar use through time. Over a long enough time, one's writing patterns may indicate a decline of cognitive ability. '[W]ith advanced statistical analysis of different writings, from blog posts to e-mails, researchers (or even employers) may pick up the downward trend of our cognitive skills long before we even suspect it' (Baker, 2007, pp. 177–178). In the near future, homes may be filled with sensors and monitors recording our behaviour (Albrecht and Mcintyre, 2005, p. 114; Baker, 2007, pp. 154–181; Tegart, 2010, pp. 11–35). Similarly, the success of Barack Obama's re-election in 2012 was credited in part to his team's use of Personal Information to reliably predict voting intention and willingness to contribute money to the campaign (Scherer, 2012). Likewise, the revelations by Snowden of national security agencies' mass collection and use of things like metadata (Greenwald, 2014) show just how interested government institutions are in our innocuous information.

These examples show that the realm of what could be classified as 'intimate information' is expanding far beyond what a patient shares with their doctor in a consultation or treatment. Moreover, intimate knowledge such as the cognitive decline of a person is drawn from fundamentally innocuous information – split second gaps between answering the phone, the form of our sentences, the way we

[9] For a recent overview of technologies associated with longevity, see: (Tegart, 2010).

[10] I discuss information in detail in Chapter 5 and in §7.4 develop an account of 'personal information' as 'information that relates to a person or group of people in some way'.

move our feet in the kitchen. The Snowden revelations underpin a larger claim that what we consider to be important Personal Information is changing in the age of surveillance, the technologies making the distinction between intimate and innocuous information dependent on the way that information is used.

What are the causal factors through which technology can change information? There are (at least) two separate, but related, factors. First is the rise of surveillance technologies and our changing behaviours. This is afforded by the 'the synergistic combination of four major "NBIC" (nano-bio-info-cogno) provinces of science and technology' (National Science Foundation and Department Of Commerce 2003). These converging technologies not only produce more relevant information about things and people but are also being developed with the capacity to *share* information across the different technological domains. Advances in nanotechnology produce information that supports cognitive technology and, when coupled with biotechnology, produce a wealth of transdisciplinary data, ready for analysis by advanced informatics (Cheshire Jnr, 2008; Hook, 2008).

The second causal factor is that this information, gleaned from a host of different disciplines, can now be collected and shared between people that were once separated from each other, either by discipline, geography, language or time (Nissenbaum, 2009, pp. 21–35). The development of surveillance technologies is astounding because of the wealth of information that it may use, the diversity of sources of this information and the incredible range of people who can access, use and ultimately benefit from this information. In the United Kingdom, the Open Health programme is a paradigm example of the ways in which Personal Information is being used for population health.[11] In parallel, another UK government agency, Government Communication Headquarters (GCHQ) has been collecting massive amounts of information. In a 2010–2011 review, GCHQ 'stated that in one 24 hour period the agency had been able to process and store "more than 39 billion events" … [meaning that] GCHQ had managed to collect 39 billion pieces of information in a single day' (Harding, 2014, p. 161). This shows just how much information we produce and how much is collected.

Focusing attention on the informational element in the age of surveillance is important, as technological and behavioural shifts mean that the standard ways that we have dealt with Personal Information may no longer be able to provide us

[11] For example, three services arising from the UK Open Health initiative and already available through the 'data.gov.uk' website, are *Dr Pocket*: 'Dr Pocket is a company that uses public information about hospitals, doctors and organisations in health to help people find the best GP for them', *The London DataStore*: 'Available for free use and reuse, however people see fit, the London Datastore has joined up with 4iP to create a development fund to encourage developers to use the raw data to develop apps, websites and mobile products' and *Health iQ*: 'Health iQ is an analytics consultancy that works across healthcare and life sciences, who helped Healthcare for London to develop a specialist stroke service.'

with clear guidance. Consider informed consent,[12] a central issue in medical bioethics.[13] If an important factor in informed consent is that a health care professional give information of any conflict of interest (Beauchamp and Childress, 2001b, p. 75), how can a remote private company accessing the UK Open Health platform meaningfully meet this requirement? These concerns become even more complex when thinking of issues arising from in-house remote monitoring, dementia and meeting the standards of patient competence necessary for informed consent to be meaningful (Beauchamp and Childress, 2001b, pp. 73–77). And in a world where the data trails we leave wherever we go and whatever we do (Nissenbaum, 2009) are being used in massive state surveillance programmes (Greenwald, 2014; Harding, 2014), concerns about Personal Information and informed consent extend to deep-seated concerns about the impact of surveillance technologies on our basic political freedoms.

Further, these records about medical visits, shopping habits, where we drive (Ramli, 2011) and who we interact with (Greenwald, 2014, pp. 160–164) are digital, so are typically[14] neither reduced by use nor limited by use-by dates: an electronic database can be accessed perpetually without any decline in the information. Assuming that the database remains stable and the technology is accessible and reliable, repeat use and access have no necessary effect on the information quality. Compare information to a pie. For each piece of pie eaten, there now remains one less piece of total pie. Likewise, as time passes, the pie gradually becomes less edible, losing flavour, nutrition, probably becoming toxic and ultimately ceasing to be food. Information in databases should not face such degradation through access or time. '[I]nformation doesn't wear out. It [can] be endlessly recycled [and] repackaged' (Drahos and Braithwaite, 2002, pp. 58–59). Given this, there is a large amount of uncertainty about what Personal Information in databases[15] may be used for, in the near and distant future. As such, those who request and provide the source information surely cannot know what they are consenting to; technologies are impacting on how we apply a principle like informed consent and can alter the fabric of our political culture.

[12] Rather than framing this as in issue of informed consent, this could be equally covered by reference to patient confidentiality. The basic issue about redundancy of key concepts in medical bioethics remains the same.

[13] Tom Beauchamp and James Childress describe informed consent as typically having seven elements; competence, voluntariness, disclosure, recommendation, understanding, decision and authorisation (Beauchamp and Childress, 2001b, p. 80).

[14] This claim is perhaps controversial, as it presumes that the technologies have stable software and hardware that do not alter the information when it is accessed and remain accessible through time. Further, as I argue in Chapter 5, *semantic* information is multirealisable, so the information can change depending on its use. However, the general claim of 'non-depletion by use' stands.

[15] This is not a new concern – people working with DNA/biobanks have had to confront it (Clayton, 2005).

The impact of these technologies on informed consent and political communities displays a larger concern about simply applying standard ethical principles to broader issues arising from surveillance technologies: the basic worry is that as the traditional patient-professional relationships and political processes break down and reconfigure themselves in the face of new technology, do we simply say that the key bioethical principles[16] are now outdated? Should we simply accept that we are under state surveillance no matter where we go or what we do? Perhaps the problem is not that the principles are wrong or outdated, but in light of technological changes, a new analysis is required of these existing ethical theories.[17]

Instead of jettisoning well-developed ethical and political principles, perhaps these changes in information technologies mean that we need to rethink[18] the moral values that underpin the principles? If so, how do we actually go about doing this? A first step in the rethinking is to be clear what we are actually talking about, 'for the way we conceptualize a problem has important ramifications for law and policy' (Solove, 2004, p. 27). These changes arise not from new moral concerns, but new ways these moral concerns are encountered in response to changes brought about by surveillance technologies. As it stands, this does not clearly describe the problem. It is not simply the convergence of the technologies and not even the informational richness coming from the technological convergence. As this book will argue, essential to understanding ethics in the age of surveillance is that the new technologies afford[19] informational aggregation, which produces an emergent Virtual Identity.[20]

This is a central claim of the book and the justifications are expanded throughout Chapters 4, 5 and 6. For now, I will simply describe what I mean by 'emergent Virtual Identity' as it relates to Personal Information. As cognitive agents, we can understand aggregated and integrated Personal Information as *an* identity: a particular identity emerges from the aggregation of Personal Information. Surveillance technologies function by bringing information together, aggregating it from a host

[16] I refer here to autonomy, non-maleficence, beneficence and justice, discussed in detail in (Beauchamp and Childress, 2001b).

[17] Unsurprisingly, this is not the first such attempt. Take Helen Nissenbaum's description of her recent research: 'The primary mission of this book is to confront and give a moral and political account of this pileup of technologies and practices, to pinpoint and understand sources of concern and to provide a framework for expressing and justifying constraints' (Nissenbaum, 2009, p. 6). My arguments, however, differ from Nissenbaum's, discussed in Chapters 2 and 8.

[18] This use of 'rethink' here is a reference to Neil Manson and Onora O'Neill's book *Rethinking Informed Consent* (Manson and O'Neill, 2007), where they argue that the model of informed consent, like that described by Beauchamp and Childress, needs to be revisited.

[19] I talk more about affordances and their special significance to identity and information in Chapter 6. In anticipation of that discussion, I will simply state here that affordances, as used in this book, relate to the ways in which technologies can make certain behaviours and/or results easier or harder.

[20] As mentioned in §1.5, throughout this book I develop a number of key terms and use them in reference to a particular set of meanings. Typically, unless otherwise mentioned, I will indicate this by use of capitalisation; Virtual Identity is one such term.

of different sources and integrating it into a Virtual Identity. 'In the Information Age, personal data is being combined to create a digital biography of us . . . In short, we are reconstituted in databases *as a digital person composed of data*' (Emphasis Mine, Solove, 2004, pp. 44, 49). While Daniel Solove and I differ in the terms we use,[21] we are both concerned about the same process, 'where a multitude of dots juxtaposed together form a picture, bits of information when aggregated paint a portrait of a person' (Solove, 2004, p. 44). The aggregation and integration of information about a person produces something new, a rich and detailed portrait of the person, a Virtual Identity. The time one gets out of bed, the way one walks into their kitchen, the time taken to recognise a voice on the phone, shopping habits, one's attraction to George Clooney: as independent data points they will tell little of interest about the person. But when large amounts of data are accumulated through time and these separate data streams are aggregated, a highly detailed and potentially intimate 'portrait' of this person emerges: being attracted to George Clooney was part of a profile used by Obama's re-election team (Scherer, 2012).

This Virtual Identity is not simply an aesthetic entity; it can be highly revealing about that person and/or can be used to harm the person. For instance, certain repeated behaviours may set off a series of triggers, indicating that the person is losing cognitive ability and may be developing dementia. However, once we consider aggregated and integrated Personal Information as a morally reactive Virtual Identity, the scope of the information that we ought to be concerned about expands dramatically: aggregated and integrated Personal Information suddenly becomes relevant far outside of the field of bioethics. We have now moved from discussing intimate information like that produced in a medical context to the moral importance of innocuous Personal Information like minor changes in speech patterns. And, assuming that we ought to treat like cases alike, given our concern about the intimacy of medical information, I suggest that we ought to be similarly concerned about the potential for innocuous information to be equally intimate. This concern extends into the fabric of our political culture. The revelation of state activity in the age of surveillance is an indication of just how important these Virtual Identities are – what's at stake is life and death and the core values of our liberal democracies.

1.3 VIRTUAL IDENTITY IN AN ETHICAL VACUUM?

Recognising parallels between medical information and that produced by surveillance technologies is important: if there's something morally important about medical information, then perhaps surveillance technologies produce similarly important information. And insofar as like cases should be treated alike, consistency may demand that we treat surveillance information similar to the ways we treat

[21] Solove favours 'digital person', while I favour 'Virtual Identity'. The explanation for this is given in §4.7.