Utrecht University

Privacy Self-Management and the Issue of Privacy Externalities

Managing One's Privacy Badly, and Others' Even Worse

Notification: your picture was automatically pixelated to respect the privacy of others. Be more careful next time.

Simeon de Brouwer

Thesis MA Applied Ethics 2018-2019

Simeon de Brouwer

Student Number 6428916

Supervisor: Dr. Joel Anderson

2nd Reader: Dr. Hanno Sauer

Utrecht University - MA Applied Ethics 2018-2019

20th of June 2019

Abstract: Privacy online is mostly a matter of self-management — or so it is expected from individual data subjects, through the 'notice-and-consent' (N&C) model of privacy-protection. However, strong evidence shows that privacy self-management is too demanding for most people. To this burden, I add the overlooked burden of respecting the privacy of others, which is threatened by the issue of privacy externalities — a natural consequence of privacy being an interpersonal matter in some respects. Regarding this new issue, I first discuss the wide range of 'negligent' practices (such as uploading pictures online or sharing one's genetic information) which generate privacy externalities (the impact on third-party subjects' privacy). I then argue that this issue cannot be solved through N&C alone without disproportionate consequences, and explore three alternative or complementary models of privacy-protection in which the tension between our values (privacy-as-control) and our convenient (and negligent) way of life may be eased. I then suggest a blend of these which, together with N&C, appears to solve both the issue of burdensomeness and that of disproportionality. The prominent element is an individual's duty of care for others' privacy, intended to internalise privacy externalities. I close by highlighting the aspects of this suggested solution for which research is most direly needed.

Keywords: interdependent privacy, externalities, self-management, duty of care, negligence, notice and consent, GDPR, accountability, burden, proportionality.

Word Count: 19,742

Source of cover's picture: unsplash.com (Creative Commons free licence)

Table of Contents

> "The law has nothing to say about whether I can prevent my sisters or cousins or my kids from putting their genomic sequence out there. […] [Some] people are giving up privacy and other people want it. The ones giving up their genomic data have the potential to infringe upon the privacy of others, but it's kind of a grand experiment. If we get it wrong now, what's the worst that will happen? We lose privacy for a generation. And then half a generation after that. And half a generation after that." The geneticists in the room laughed. The lawyers, not so much.
>
> Molteni 2019

1. Introduction

Privacy has become in the past decade a much-debated topic in legal, societal and philosophical circles, and has gained even more attention through its most recent upgrade, European Union's (EU) the Regulation 2016/679, or General Data Protection Regulation (henceforth GDPR). While this upgrade addressed mostly the aspects of privacy-protection related to its enforcement, the GDPR's essential method of privacy-protection remains unchanged despite having been widely criticised in the past. This method is the provision of rights to self-management, enforced through processes of Notice-and-Consent (hereafter 'N&C').

The degree of attention and the number of tasks that individuals are expected to perform with regards to their privacy is already so burdensome that most people do not manage their privacy much — if at all. The GDPR addresses this issue in various ways, as I will show. Yet, those who do not care about their own privacy may care even less about that of others, and I will show that, through everyday behaviour, one can remarkably impact the privacy of others.

For example, we routinely upload pictures of others to proprietary platforms such as Facebook, even though we know these are not necessarily safe places to display or store data about our private life. We disclose the genetic data of our whole family, together with ours, when we get DNA testing kits from companies such as 23andMe. The discussions we have with our visiting friends fuel the training of Amazon's AI when they enter our Alexa-equipped 'smart home.' None of the aforementioned individuals benefit from adequate privacy protection: the means we rely on to ensure privacy protection, such as contract-like Terms of Service (TOS) between user and service-provider, provide N&C to the user only. This is problematic, if self-management of one's privacy is supposed to be attainable (instead of an ideal to strive for).

This new concern is moreover different from more regular issues of disclosure such as leaks and hacks: it is not only about *bad* privacy management, but also about *impossible* privacy management. Privacy cannot adequately be managed alone, as it is in some aspects necessarily an interdependent matter, whereby privacy externalities are imposed onto others. We keep damaging others' privacy, without even being aware of it most of the time, and without being in any way held accountable for it. I will frame this as an issue of negligence — i.e. of failing to perform one's duty of care for others — but will first argue that it cannot easily be addressed through N&C alone. My research question itself is "*To what extent can the N&C model of privacy-protection address, unaided, the issue of privacy externalities, and what could (or should) supplement it in this task?"*

I will argue that, if we are to remain committed to N&C, and are looking for an easy solution to its limitations, this issue could be addressed with the development of various technological tools that will ease the (already) excessive burden of privacy management. Yet, because these tools may never come — or if they do, they may come at too high a price — we may have to change either one of two things. We may have to abandon (or soften) our commitment to N&C and to privacy-as-control (self-management), or we may have to abandon (or adapt) certain of these practices of ours which affect the privacy of others. To give a better appreciation of these possible strategies for the future of privacy-protection, I engage in thought-experiments which allow me to experimentally explore one scenario for each kind of strategy.

While each of these strategies will represent concrete alternatives or complements to our current way of protecting privacy, the value of examining them will reside more in the exploration and illustration of possible directions to follow. The goal is to indicate possible paths for future research to explore, thereby questioning our commitments and our expectations. The solution I will finally suggest will hold on to N&C, while incorporating elements of each of these aforementioned three strategies: it will address the burdensomeness of N&C with elements mostly from strategy 1 and 2, and by coupling N&C with strategy 3's duty of care, it will solve the issue of privacy externalities without having the disproportionate effects that relying on N&C unaided would have.

The method employed to answer my research question therefore consists in analysing the current paradigm of privacy-protection, its development and its known shortcomings; doing research on cases of interdependent privacy to show the magnitude of the issue of privacy externalities; and exploring, through thought-experiments, strategies that solve the tension

between our values (privacy-as-control) and our practices (where one's privacy depends on multiple people's good will).

The argument and structure of this thesis go as follows:

1. The Notice-and-Consent model, a form of privacy self-management, is the current standard of privacy-protection. The GDPR is the latest implementation of this model in the EU. Individual freedom (insofar as it is instantiated by the concept of privacy-as-control) is the core value of the model. (Chapter 2)
2. There are known issues with the N&C model (burden of self-management), which have at least partially been addressed in the GDPR. Another important challenge to N&C has hitherto not been properly addressed, however. This challenge is the interdependent aspect of privacy, where negligence generates privacy externalities. (Chapter 3)
3. It would prove to be excessively burdensome and disproportionate to persevere with N&C alone in the case of privacy externalities. Three strategies alternative or complementary to N&C are thus considered. (Chapter 4) Each on its own is problematic, but I suggest that the appropriate solution is found by blending them. (Chapter 5)
4. The concept of privacy-as-control, the expectation of self-sufficient privacy management, and the way these are implemented through a model of privacy-protection that mostly relies on N&C, should be thoroughly questioned. Under this light, the very idea and expectation of self-management become problematic, as they overlook an aspect of the individual's privacy that necessarily depends on others, and which allows for the creation of privacy externalities. Privacy self-management as it currently is implemented does not, and cannot alone, provide everyone with meaningful control over their data. Therefore, alternative or complementary modes of privacy-protection should be sought. (Conclusion)

2. Notice-and-Consent as Privacy Self-Management

The concept of privacy has been receiving an increasing degree of attention from industry, regulators and the public in the past years, especially in the Western World. The "digital turn" has meant that living some sort of 'digital life' is no longer optional, as the edges of the 'online' world are less distinguishable as it overlaps with the 'offline' or 'real' world (Bernal 2014:ix; EDPS Ethics Advisory Group 2018), and as most activities now leave digital trails. It has especially meant the increased datafication[1] and digitization of our lives, greater ease of sharing data over the internet, platformisation of services, and the growing power of large corpora (EDPS Ethics Advisory Board 2018:11, 15-16). In the face of these phenomena, the calls for renewed protection of powerless citizens led in the EU to the GDPR (Pascalev 2017:40; Floridi 2018:6).

The GDPR is the current state-of-the-art EU legal guidance with regards to individuals' rights to data protection, which in turn are what allows for privacy-protection. Besides adding significant emphasis on the enforcement of these rights however, the GDPR's strategy to protect privacy remains very much in line with that of the preceding fifty years of regulation; this strategy is the individual's self-management of her personal information (Obar and Oeldorf-Hirsch 2018:2).

Information-privacy protection has, since its inception, largely been framed as a form of giving control back to the individual over how she makes herself known to the world. The modern concept of privacy began with an article by Samuel Warren and Louis Brandeis, "The Right to Privacy," which argued for the recognition of a new legal right to "be let alone" (Warren and Brandeis 1890). The article was spurred on by the increasing use of the photograph, which was felt to be unduly invasive. Relying on legal analogies (especially slander and libel), the authors noted that

> The right to privacy does not prohibit any publication of matter which is of public or general interest. […] The right to privacy ceases upon the publication of the facts by the individual, or with his consent. […] The truth of the matter published does not afford a defence. […] The absence of 'malice' in the publisher does not afford a defence. (Warren and Brandeis 1890:214, 217-8)

Warren's and Brandeis' conception of the right to privacy thereby concerned only very 'personal' information, in the sense of information which has no reason to be known publicly,[2]

[1] The transformation of all kinds of information into machine readable, mergeable and linkable form (Taylor *et al.* 2017:3).
[2] Such as emotions, whether expressed in conversation or in facial expression (Warren and Brandeis 1890:206).

except with the consent of the individual in question.[3] This early focus on consent remained intact, and even increased, throughout the development of (Western) privacy-protection regulations (such as the 1973 (US) Fair Information Practices Principles,[4] the 1980 OECD Privacy Guidelines or the 1995 (EU) Directive 95/46/EC) and is the reason that privacy is currently considered to be mostly a matter of self-management.

Due to the importance ascribed to consent, privacy management often takes the contractual form of two parties agreeing about the processing (i.e. collection, use, disclosure, etc.) of the 'user's' personal information, in exchange for a service offered by the 'provider.' This contractual form is no accident: the protection of privacy I am focusing on here is sought and developed in the West (EU – US), where the prevailing political doctrine in the past century has been liberal capitalism (Cf. Mindle 1989). In this framework, primacy is given to the individual,[5] and (paternalistic) intervention by the state in private matters is limited. The individual's consent to a contract's terms is under these circumstances most often sufficient to legitimise the agreement, as long as the consent is informed and voluntary — even if there are certain setbacks to the interests of one of the parties.

This can be seen through the importance ascribed to freedom (conceived as self-determination, or control)[6] in Western civilisation — and especially in the EU, on which I will focus. Individual freedom is one of the central concepts of ground-setting documents crucial in European law and traditions, such as in the Charter of Fundamental Rights of the European Union and in the 2007 Treaty of Lisbon .

This importance of freedom led to the definition of privacy as the relative control over the ways and the extent to which one selectively discloses (information about) oneself to others.[7] Information privacy, ensured through specific rights to data protection, is therefore a

[3] In contrast, 'personal' information following the GDPR refers to any information about an identifiable person (GDPR art. 4.1).
[4] See Gellman 2014.
[5] This is also largely due to the regulation appearing the wake of early data analytics, the scale of which was largely constrained by its limited capabilities.
[6] For the sake of brevity, I will bracket questions about the distinct conceptions of freedoms (negative, positive, and Republican) and simply focus on 'freedom as control' because of its importance in what follows.
[7] I say 'relative' here, because there are certain degrees of control which fall under the concept of 'having privacy.' It is difficult to specify what degree of control is required, especially as privacy is at least partly subjective and context-sensitive. Cf. Kupfer 1987.
For references to this understanding (i.e. definition) of information privacy, see Westin 1967:7; Culnan and Armstrong 1999; Culnan 2000; Weinreb 2000; Hann *et al.* 2002; Whitman 2004:1161; Alge *et al.* 2006; Moore 2007; De Hert 2008; Whitley 2009. Although privacy may be defined in various ways (Introna 1999; Solove 2002), the definition I am here focusing on is the one that remains the most influential today — especially considering how privacy is protected by the GDPR, which focuses on personal data protection through individual rights of control. The reader might disagree with this and argue that the common denominator of privacy is rather something like secrecy or intimacy (cf. Solove 2002), but this does not affect my thesis, insofar as 'control' remains an important part of the concept.

matter of control, and the expectation is that individuals will manage their privacy themselves (Whitley 2009; Reidenberg *et al.* 2014; Taylor *et al.* 2017:6).

Privacy self-management is pursued through a specific framework: what I call the Notice-and-Consent model ('N&C').[8] This framework sets out standard procedures which ensure uniformity and consistency in privacy self-management, which in turn ensures that individuals have the means to meaningfully manage their privacy themselves. As its name indicates, the N&C model requires that data subjects are notified of the intended processing of their data by the data controller,[9] and that they consent to it. The aim is that the subjects be aware of the transaction (their data in exchange for a service), and that their consent be informed — i.e. a voluntary and non-deceitful contract between free and autonomous parties.

Although there can be additions (e.g. restrictions regarding which uses of data are allowed in the first place) and exemptions (e.g. data processing for journalistic purposes) to the decentralised framework of N&C, it can adequately be termed the central pillar of privacy-protection — even in the wake of the GDPR, which brought to the fore other criteria for privacy-protection, such as privacy-by-design requirements and the principle of fairness in the use of data.[10] Control is assured on a one-to-one basis in the form of a contract between a user and a data controller, whereby the user agrees to certain uses of her data.

I will now argue that, although the GDPR was an attempt to empower European individuals (EDPS Ethics Advisory Board 2018:6, 15-16; Floridi 2018:6), its substantive reliance on the N&C model of privacy-protection is self-defeating. Simply put, asking someone to do more than she possibly can, will lead to either under-fulfilment or poor performance. The main contribution of this thesis, however, resides in showing that even more needs to be done.

[8] Cf. Hull 2015. It is sometimes called the 'notice-and-choice' or 'informed consent' model.
[9] A data controller determines the aims and the means of processing personal data. It does not process *per se* the data itself, as this is done by the data processor; however, for simplicity I will discuss data processing as if carried out by the data controller.
[10] Cf. FTC 1998:7; Koops 2014; Hull 2015.

3. The Shortcomings of the Notice-and-Consent Model

In this chapter, I argue that managing privacy adequately, in the sense of achieving comprehensive, sufficient protection, is extremely difficult, and is thus in practice not a right available to everyone. There are several issues preventing individuals from achieving this goal of meaningful control in privacy management. The most prevalent concerns the user's excessive (cognitive) burden. Another, hitherto mostly overlooked in the debate, concerns the (lack of) control of those who are mere third parties to the user's behaviour, which creates negative externalities (a loss of privacy for these third-party subjects).

Policy-makers are aware of some of these issues,[11] and modern regulation attempts to address the shortcomings of the N&C by adding other N&C-related clauses, as I will show. However, I will argue that the consequences of interdependent privacy cannot viably be addressed in this way, and that this has implications for our ideals of self-governance. The goal of this section is to introduce the shortcomings of the N&C model most relevant to this thesis. I begin with known criticism faced by N&C and show how the GDPR addresses it, after which I move on to the less familiar issue of interdependent externalities that is the focus of this thesis.

3.1. The Over-Burdened User

A privacy policy is usually not user-friendly and accessible; it was, after all, a document created by lawyers for lawyers, not a consumer tool (Litman-Navarro 2019). The language used may not be accessible to everyone because the policies sometimes describe complex mechanisms, use legalistic or technology jargon, and even purposefully obfuscate the meaning of certain parts to make them cover more ground (Anton *et al.* 2004; Nissenbaum 2011).

Even if they had the knowledge and skills to do it, however, reading TOS and privacy policies[12] would simply take too much time for most individuals (McRobb and Rogerson 2004; Milne and Culnan 2004; Whitley 2009). In the year 2008 alone, it would have taken 180-300 hours for an average US citizen to read the privacy policy of every website they visited (McDonald and Cranor 2008). This would have meant for the US a 'loss' of about 40-67 billion hours of productive time per year, spent on reading these documents instead of doing something else, estimated at 600-1100 billion dollars' worth of productive time. In 2019 these numbers are very likely to be much higher, due to the increasing digitisation and platformisation of

[11] For example, see FTC 2012; OPC 2017.
[12] Since both privacy policies and TOS need to be read for having a meaningful grasp of the data the user allows to be processed, I will henceforth also implicitly refer to the TOS when mentioning privacy policies.

services (i.e. non-website-based services and everyday 'smart' objects which collect their user's data would be relevant to this calculation, in addition to the websites considered in the 2008 calculus).

Not only are most people unwilling to spend that much time and energy on privacy policies; even more people simply cannot do so, because they are already struggling to make ends meet with the time and energy they have. As argued by some (Newman 2015; Madden *et al.* 2017; Madden 2019), the fact that those with lower incomes are less likely to take privacy-protection measures when online places these people at greater risk of online targeting and exploitation. This raises concerns about how in this age of big data, information inequality is transferred into economic inequality (Trappel (ed.) 2019).

These are serious obstacles to the ideal of privacy self-management, if privacy is a (human) right that everyone should be able to enjoy. The situation was far worse a decade ago, however, as it was common practice for data controllers to provide little in the way of N&C. Privacy policies had too broad a scope to be in any way meaningful, and consent was either assumed[13] or binary,[14] meaning no real control was given to the user. Moreover, the privacy policy of the data controller was not the same as that of the third parties with which users' personal data was shared, and all these policies could change overnight without notice, which meant users had to regularly consult many privacy policies if they wanted to provide actual informed consent.

The GDPR addressed these issues by strengthening the requirements of N&C and increasing the individual's control over her personal data. This involved requiring privacy policies to be as short, clear and concise as possible. It also involved providing a strict definition of informed consent by which to abide, and requiring that the particular permission given by the user regarding the processing of her data are applied throughout the chain of third-party data controllers accessing said data. Notably, the GDPR made the compliance to these conditions a reality by designating accountable agents in each step and in every organisation processing large quantities of data, by establishing mechanisms and bodies of oversight, and importantly, by stipulating a large fine for non-compliance (art. 37 and 83).

Thereby, the GDPR made sure that the agreement between user and data controller would fulfil the modern contractual standards, i.e. that the contract involves no deception or

[13] That is, using the services was taken to imply the user consented to the terms.
[14] That is, either completely allowing the data processing, or completely rejecting it. If chosen, the latter was not even always taken in consideration if the user still used the services. Cf. Strahilevitz 2010:2038.

forced choice, that both parties can choose (to a certain extent) their own terms,[15] that the contract applies throughout the 'life-span' of the data, that there are certain penalties for breach-of-contract, etc. With these requirements, the GDPR greatly simplified and lightened the burden of privacy self-management. "With each sign of failure of privacy self-management," says Solove (2013:1882), "the typical response by policymakers, scholars, and others is to call for more and improved privacy self-management."[16] While this is not completely true with the GDPR, privacy self-management remains the crux of its privacy-protection strategy.

All there remains now is that the user does her part. Still, despite these improvements, self-management remains marginal. In theory, people are given the means to achieve some form of privacy self-management. In practice however, most people do not actually manage their privacy, because it remains too burdensome: neither do they bother reading privacy policies, nor do they opt-out when given the opportunity to do so, nor do they change default settings (Solove 2013:1884). The improvements when privacy policies are expressed in clear tables, icons, labels or short texts lead to only marginal improvements in people's understanding and behaviour (Calo 2012:87). Moreover, in 2019 most policies remain incomprehensible to the average user (Litman-Navarro 2019) and consent acquisition remains suboptimal or deceptive (Privacy International 2019a).[17]

The GDPR's new requirement of 'privacy by design and by default,' though a safeguard for individuals' privacy, does not raise the proportion of privacy self-managers. Instead, it might even lower it, since people can now rely on the expectation that the standard is the least privacy-invasive — hence that they already are 'safe.' Participation in the N&C model (clicking "I have read and agree to the conditions") remains "the biggest lie on the internet" even today, despite the additional safeguards and improvements (Obar and Oeldorf-Hirsch 2018).

This altogether leads us to question the self-management model of privacy-protection, and the ideal of liberal (freedom conceived as) control behind it. The model of N&C was built when the capture and use of data was considerably less extensive and complex than today (Taylor 2017:6, 8), and might therefore be somewhat obsolete. While this model imposes an important burden on individuals to manage their own privacy, however, I will now show that consistently applying it to the less familiar issue of privacy externalities requires even more efforts from data subjects.

[15] E.g. when choosing between a certain range of cookies, such as those strictly necessary for the service to be possible and those that allow certain forms of marketing and advertisement.
[16] Cf. Nissenbaum 2011.
[17] See also 'dark patterns' (Brignull 2019).

3.2. Interdependent Privacy and Privacy Externalities

Thanks to the GDPR, one has various rights with regards to one's personal data held by a data controller. Among others, one has the right to access this data. Therefore, data controllers like Facebook have a portal where users can access their data; yet, this access seems to be restricted to "information you've entered, uploaded or shared" and "information [we] associated with your […] account."[18] This restricted access does not fully reflect the broad scope of GDPR-endowed rights to data protection, the definition of personal information of which is

> any information relating to an identified or identifiable natural person ('data subject'); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person. (GDRP art. 4.1)

This definition is purposefully extensive, so as to protect data subjects whenever information about them is processed, not merely with regards to information they provided themselves directly. However, if this is the definition which establishes the boundaries of 'personal data,' it seems that certain categories of personal data are not (fully) taken into account by data controllers, such as in the Facebook example mentioned above.

Let us look, take at the various distinctions made by Bruce Schneiders (2010) between different types of 'user' (i.e. personal) data for Social Network Services (SNS), which apply to most other data-related services in order to further illustrate this:

- Service data is the data you give to a social networking site in order to use it, e.g. name and address.
- Disclosed data is what you post on your own pages.
- Entrusted data is what you post on other people's pages.
- Incidental data is what other people post about you.
- Behavioural data is data the site collects about your habits by recording what you do and who you do it with.
- Derived data is data about you that is derived from all the other data, e.g. inferences.

[18] Facebook 2019 (https://www.facebook.com/your_information/, accessed 02/06/2019).

While this taxonomy could be questioned — as other ways of looking at personal data may also be adequate[19] — it includes an aspect of privacy still rarely taken into account by data controllers. This aspect is the role that "other people" (henceforth 'third-party subjects,' as opposed to the 'user' of the service who discloses their personal data) play in the processing of one's personal data by a data controller.

Indeed, as much as posting content about myself adds to my personal data which is processed by the data controller, the content about me posted by others also adds to it. When I upload material on Facebook, it is related to me (and therefore is my personal data) in that it is uploaded *by* me, and *about* me —two relations of 'identifiable relatedness' relevant for constituting personal data, together or individually. Accordingly, when someone else uploads content about me, it is both their personal data *and* mine — as long as I am identifiable — because, although it is uploaded *by* them, it is clearly *about* me. I will call these relations 'causal agency' and 'personal relevance.' An instance of the 'causal agency' relation is that a data subject is tied to a certain content in virtue of her user details (username, browser identity, etc.) being linked to the action that made the content available. An instance of the 'personal relevance' relation is that a data subject is identifiable insofar as the content itself contains or *is* information about her which can be used for the identification process — such as a picture.

This distinction is one of my main contributions, as it allows to identify and frame a major obstacle to privacy self-management that has been hitherto largely overlooked by the general public, data controllers and policy-makers. This obstacle is the (necessarily) interdependent nature of some aspects of privacy, which in our case is coupled with the (contingent) lack of individual accountability for the impact one has on others' privacy. It is an obstacle insofar as privacy is framed as an individualistic matter, an aspect of one's life which is under one's control (i.e. self-sufficiently manageable). This obstacle I raise has two sides: one is enforcement-related, the other is more structural. The first is that your personal data is 'your personal data' independently from how it was arrived at, and that it should accordingly always benefit from the protection framework afforded by the GDPR. The fact that this protection has not consistently been distributed to all instances of personal data is only an issue of awareness and compliance (and thus of enforcement). The second side is that the behaviour of a user sometimes generates privacy externalities, unfavourable to others.

An externality occurs when the price paid for a good or service does not reflect its full costs. In the context of privacy, people's decisions to share their personal information may

[19] See for instance the other categories offered by Schneiders 2010 or by Kitchin 2014b:4. These do not address (as explicitly) the role 'others' play in the data controllers' processing of one's personal data.

allow the data controller to know more or better about others. People thus think they 'pay' for the service with their data, but some of that data is actually (also) other people's data. The full costs of the service include the impact on others' privacy and the (dis)utility resulting therefrom, and this is neither transparent nor accounted for in the transaction between user and service-provider. Hence the term 'privacy externalities' (Laudon 1996:14-6; Humbert *et al.* 2015; Symeonidis *et al.* 2016:2; Choi *et al.* 2019:7).[20]

The issue of non-compliance is a serious one, and I will address it by arguing below that a broader range of stakeholders (than just the user) should be involved in the N&C process of certain data practices. However, it is the issue of privacy externalities that will be the main issue here, as it brings the concepts of negligence and accountability into the picture, and is thus more interesting ethically-speaking than an issue of mere compliance. In an age where an increasing proportion of our lives is digitized and recorded, it is crucial to examine how much of others' acts, lifestyles and thoughts we capture or help to disclose, inadvertently and unknowingly. Despite all efforts one makes to protect it, one's privacy can still be breached because of the behaviour of others, and as danah boyd (2011) succintly puts it, "[e]verything that everyone else does that concerns you, implicates you, or might influence you, will go down on your permanent record."

Disclosure of others' personal data through one's activities can be repetitive, commonplace, extensive and substantial, even though it happens mostly because of negligence: it is thus a serious issue. If we want to be consistent in the way in which we protect people's privacy, (the current standards of) privacy-protection should apply in these 'new' cases of personal data disclosure too. I will show below how demanding that could be; for now, let us have a closer look at cases of interdependent privacy — of which I distinguish four categories — and especially those of which are morally problematic. These categories are:

1) Direct disclosure: data is revealed about subject 1 when subject 2 discloses data about subject 1.
2) Indiscriminate sensing: data is revealed about subject 1 when subject 2 reveals data about subject 2 that was formed through an indiscriminate process of capture, and which therefore included data about subject 1 alongside the data of subject 2.
3) Fundamentally interpersonal data: data is revealed about subject 1 when subject 2 reveals data about subject 2, which necessarily is also data about subject 1.

[20] See also MacCarthy's (2011:24-5) narrower account of privacy externalities).

4) Profiling: data is revealed about subject 1 even though no one discloses data about subject 1. Subject 2 discloses data about subject 2, which others use to infer data about subject 1.

The difference between categories (2) and (3) is that in the former, the interpersonal data (i.e. data being about more than one subject) is only contingently so, whereas in the latter it is necessarily so. In the former, the data could have been only about the user, if she had been cautious; that is not an option in the latter category. The distinction will become clearer with examples from each category:

1) As long as it is data-recorded,[21] any activity that consists in explicitly talking *about* someone counts as revealing that person's personal data, and thus as an activity relevant to interpersonal privacy. This includes blogging (Solove 2007:24) and giving nicknames to people (Privacy International 2019b, about the TrueCaller app). This category also involves directly handing over other people's data, like when Facebook apps ask the user to access her friends' list and their data (Besmer and Lipford 2010; Hull *et al.* 2011; Bizcok and Chia 2013; Symeonidis *et al.* 2016).
2) Recording one's voice or environment often also implies indiscriminately recording others. Sensors capture all the available data of a given category (say, sound and/or image), and do not discriminate between user and non-user. Therefore, taking a selfie in a crowded place (see the case of pictures below) or installing a 'smart assistant' in one's home will also capture the personal data of other people — strangers, friends, relatives, visitors, etc. — who may neither be aware nor capable of resisting the invasion of their privacy. Recording events (in sound or image) can be a sensitive practice, because many personal aspects of one's and others' life can be thus made available to data controllers, including sensitive data like political opinions, religious beliefs, or health data (Vallet and Bonastre 2017; Vallet 2019). This data can moreover be automatically 'mined' by image-processing, voice-processing, and facial-recognition AI. This category is quite broad, but besides mere pictures and voice-

[21] What differentiates 'privacy' from 'information privacy' resides mainly in the digital world's characteristics of borderless ease of access, of reproducibility, of sharing, and of processing of information, as well as in the near-infinite life of the information while in digital form (i.e. data). The information at stake can be the same in digital and non-digital cases, but in the context of information privacy, it is infinitely more potent because of these characteristics, and because of the insights that can be drawn from them.

recordings it also includes cases such as CCTV (ICO 2017), IoT (Livingstone 2013), smart homes and smart cities (Kitchin 2014b; Quain 2018).[22]

3) There are some kinds of data which necessarily constitute personal data of multiple persons. This includes relational data (Jernigan and Mistree 2009; boyd 2011; Backstrom and Kleinberg 2014) and more particularly address-book sharing (see the case of contact lists case below). Also included are data from smart grids (McDaniel and McLaughlin 2009), data about groups (such as households or neighbourhoods) (Taylor *et al.* (eds.) 2017). Importantly, genetic data is also concerned: giving rights to a data controller to process your genetic/genomic data not only affects you and your privacy, but also potentially countless[23] individuals to whom you are related — knowingly or unknowingly (Chadwick *et al.* (eds.) 2014; De Hert 2017; Hallinan & De Hert 2017; Taylor *et al.* 2017:9; Erlich *et al.* 2018; Smit 2018; Hallinan & Molteni 2019).[24] Because certain genetic traits are necessarily shared with family members, it suffices that a single person undertakes such an analysis for a kind of 'family-wide sharing of personal data.'

4) When enough people disclose ample information about themselves, data controllers (or data brokers) are able to understand the relation between having a given trait and a specific characteristic. For example, there is a correlation between, one the one side, buying felt-pads to prevent one's furniture from scratching the floor, and on the other side paying one's bills on time (Duhigg 2009). When correlations like these have been found (through mining the massive troves of data made available to data controllers and brokers by willing individuals), the small, seemingly-insignificant pieces of personal data that prudent people disclose (willingly or not) will reveal more data about them, whether they like it or not (Choi *et al.* 2019:8, Wachter and Mittelstadt 2019). This is the case of the 'made-up' groups from profiling categories and algorithmic data-mining. Another type of case is that, when a certain practice is beneficial in a given context and widely adopted, failing to disclose whether or not one does that practice can be taken to imply that one has "something to hide" (i.e. pertains to the category of those who would suffer if the answer was known). For instance, this happens when insurers ask about smoking habits (as it is assumed that only non-smokers would have personal

[22] See also other 'smart' environments such offices (Amazon 2017), hotels (Fox 2018), rental properties, stadia, hospitals, and senior homes (Mims 2019).

[23] The question is open whether dead people and people yet to be born have a right to (genetic) privacy, but regarding the living that should be uncontroversial.

[24] Cf. Tavani 2004. Even though the use of genetic data is highly restricted in Europe, genetic data is so rich and potent that Europeans could be indirectly affected by what their distant US relatives disclose about their own genetic information (Endedijk and van den Berg 2019). The accessibility of genetic data is already somewhat problematic (cf. Erlich *et al.* 2018; Brown 2019; Aldhous 2019), and with leaks or hacks this could also worsen.

incentives to disclose this information) or driving habits (only bad drivers have incentives not to get a telematics device). Thus, the very non-disclosure of data can already be a data point — one revealing negative traits — because of the behaviour of others.[25]

These categories (which may overlap) and examples show how important and diverse the cases are where one's behaviour can damage the privacy of others, and thus that the issue at stake here is not a rare or minor one. In some cases, such as with biometric data, the data can be very sensitive, and the impact of the disclosure can be lifelong. Negligence from the 'user' who discloses others' data may, in these cases, significantly impact the privacy of many people, but is foreseeable, and should be internalised. Moreover, even the smallest disclosures are not insignificant, due to the possibility of the data to be sold to (and aggregated by) data brokers — third parties whose business is then to exploit the data (Symeonidis *et al.* 2016; Choi *et al.* 2019:8).[26]

Issues originating from improperly-addressed cases of interdependent privacy have been directly[27] addressed in the following terms by the following authors:

- 'collateral damage' and 'spillover' (Hull *et al.* 2011; Symeonidis *et al.* 2016;[28] Tucker 2017);
- 'interpersonal management of disclosure' (Lampinen *et al.* 2011);
- 'networked privacy' (danah boyd 2011; Lampinen 2011; Marwick and boyd 2014);
- 'interdependent privacy' (Biczok and Chia 2013; Symeonidis *et al.* 2016);
- 'privacy leak factor,' 'shadow profiles' and 'online privacy as a collective phenomenon' (Sarigol *et al.* 2014);
- 'privacy externalities' (Laudon 1996:14-6; MacCarthy 2011; Humbert *et al.* 2015; Symeonidis *et al.* 2016:2; Choi *et al.* 2019);
- 'genetic groups' (Hallinan & De Hert 2017).

In most cases, the authors only addressed the problematic phenomenon in relation to either social media or genetic data, and because they addressed it in that specific context, the analysis

[25] Here, N&C functions perfectly (as information is shared only with consent) but is not enough to prevent the acquisition of data, because the very act of (non-)consenting is revealing (cf. MacCarthy 2011:26-7).
[26] Data brokers are often the main revenue of free service besides advertising, because they buy users' data *en masse*. The user agrees to this by giving her consent for her data to be shared with the data controller's "partners."
[27] This issue has moreover been more tangentially or briefly touched upon in the following sources: Bloustein 1978; Roessler and Mokrosinska 2013 (the 'network effect'); Kitchin 2014a ('data shadows'); Hull 2015:97; Taylor *et al.* (eds.) 2017 (some aspects of 'group privacy'); Facebook Inc. 2018 (sharing of one's friends information with third-party apps in the Cambridge Analytica scandal); Garcia-Murillo and MacInnes 2018.
[28] Symeonidis *et al.* conducted a survey demonstrating the concern of individuals with 'collateral damage' (externalities). They and Choi *et al.* (2019) moreover computed the likelihood and significance of the phenomenon (in the context of Facebook third-party apps or websites).

and solutions they provided are also mostly restricted to such contexts. In contrast, I argue that the scope of the problems resulting from improperly-addressed interdependent privacy is much broader and therefore much more serious than it appears from the existing literature: not just a localised problematic practice, but a common phenomenon that threatens the ideal of privacy self-management itself.

I will now analyse in further detail two cases of categories (2-3) distinguished above: contact lists (fundamentally interpersonal data) and pictures (indiscriminate sensing). These categories are particularly relevant, because they are more than just cases of interdependent privacy: they are cases of negligence too — a kind of negligence for which one can reasonably be held accountable — whereby externalities are generated. In contrast, the wrongness of category (1) rests less in foreseeable violation of a certain duty of care (the corollary of negligence) than of a duty not to harm, since the disclosure is very direct and cannot really be framed along then narrative of negligence and externalities (or not as neatly as the two subsequent categories). Moreover, privacy externalities from category (4) are less foreseeable by the individual and cannot be avoided her alone.

3.2.1. Contact Lists

Contact information is often shared by users of messaging services, such as the smartphone app Facebook Messenger (see figure 1).[29]

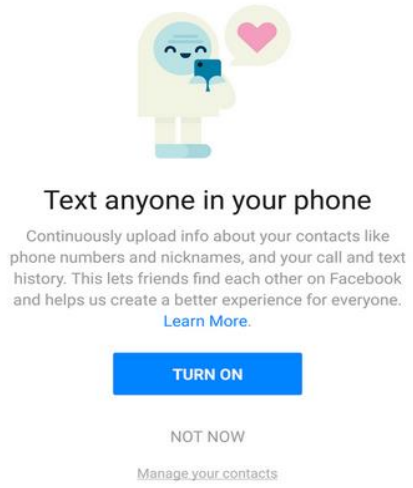

Figure 1: Facebook Messenger's Permission Request to Access the User's Contacts

[29] People do not always have the choice to even choose to manage their own privacy (and that of others), however: in 2019 Facebook made the upload of (email) contacts compulsory for users signing up with certain email providers (EFF 2019).

While it is advertised as important for the app to function well, this practice of sharing contact information has received recent criticism, as it was shown that Facebook uses this data for other purposes, more particularly for advertising (Hill 2017; Gebhart 2018; Hill 2018; Venkatadri *et al*. 2019). This case might involve deceptive N&C,[30] but not only. Contact information is very personal (data): it may contain names, nicknames, private and work addresses, birthday date, social profile links, profession, call history, etc. It would thus seem that, under the GDPR, Facebook should notify the individuals to whom the data is 'personally relevant' (i.e. the contacts), and ask their consent for its processing. The user's consent alone is not sufficient to process her contacts' data.

Even if only someone's nickname and a phone number were disclosed, it would remain a potentially fruitful acquisition for the data controller. Indeed, phone numbers are relatively reliable identifiers across services, as most people have only one (or two) and tend not to change it (them), for convenience.[31] By comparing this data with other datasets — through the services of data brokers, for example — it would be relatively easy and cheap for a data controller to gain deep insights into the lives of the many people whose data was disclosed by the user and her contact list. This, in turn, could for instance be exploited for targeted advertising purposes.

Moreover, even if no additional data is gained through other datasets, the widespread disclosure by individuals of their contact list would already be significant for the privacy of the

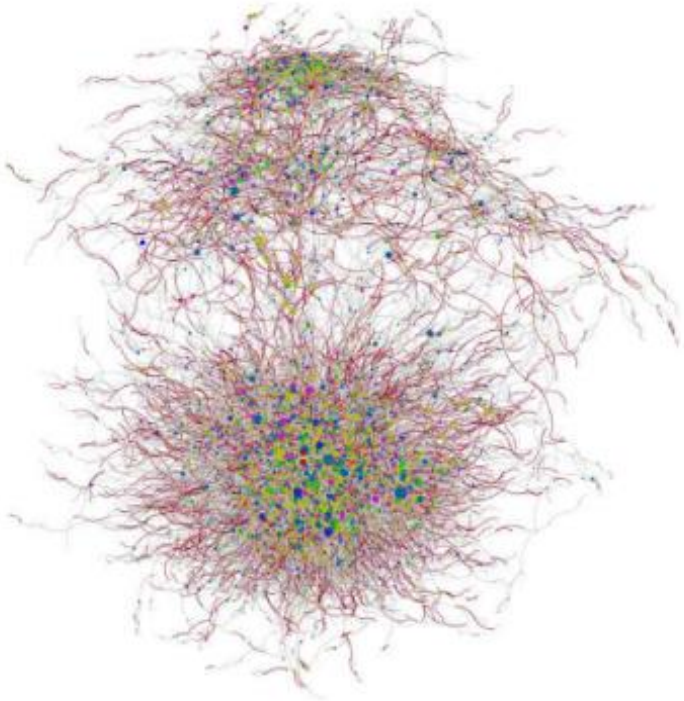

Figure 2: The network for a subset of Friendster users

[30] As well as what Innes (2001) calls "control creep," repurposing data originally gathered for a different goal.
[31] All contacts would have to update the person's phone number, were she to take a new one.

individuals concerned. Imagine I am extremely concerned about Facebook and its activities, and do my best to remain "under the radar." If, for example, five people who have my number opted to upload their contacts to Messenger, this would be enough for Facebook to identify the overlap and begin the creation of a file about me, the first data points of which would be my phone number and the relation I have with identified people — visualised as a network map such as in figure 2 (Sarigol *et al.* 2014), where this was done experimentally with the data available from Friendster, a pre-Facebook social network.[32]

Even solely knowing about this network of relations is valuable to the data controller, following the principle of homophily. Homophily is the principle according to which people are likely to interact with others who are similar to them, which means that from people's communication networks we can identify their contacts' "ethnicity, gender, income, political views and more" — all categories of sensitive data (Caughlin *et al.* 2013:1). Empirical studies by Sarigol *et al.* (2014),[33] as well as the 'Gaydar' experiment (Jernigan and Mistree 2009), have shown how effective this can be in inferring sexual orientation — a type of personal data particularly sensitive in certain parts of the world. Thus, my ability to remain under Facebook's radar is heavily undermined by other individuals' seemingly innocuous actions, which not only disclose information about them, but also (foreseeably) about me — even if I am not a Facebook user myself.[34] This is where the expectation of privacy self-management model becomes problematic again.

3.2.2. Photos and Facial Recognition

The second case examined in detail concerns photography. Taking pictures is nowadays mostly an innocent and common act. However, this act and the practices that go with it (sharing, storing, editing, etc. the pictures taken) should be made to conform to privacy regulation, especially in light of technological evolution.

As mentioned in chapter 2, the advent of photography was also the catalyst for the emergence of the modern concept of privacy. Photography was seen as a threat, because it allowed for the unidirectional capture of moments (meaning the action of one is enough to implicate others), private or public, to which boundaries needed to be set to avoid that "what

[32] The phenomenon represented through this network map is 'modularity,' and has widespread implications in many fields — including privacy protection. See for example Caughlin *et al.* 2013.
[33] Another study by Garcia (2017) and supervised by Sarigol reproduced the results.
[34] This was discussed is the notorious case of 'shadow profiles.' See Schrems 2011, Knibbs 2013, Sarigol *et al.* 2014, Garcia 2017, Hill 2017.

is whispered in the closet shall be proclaimed from the house-tops" (Warren and Brandeis 1890:195).

While the concept of privacy has evolved greatly since then, so has photography. Nowadays, cameras can have high resolution, be as small as a grain of salt, film at night or zoom kilometres and still yield great quality pictures. Moreover, their relatively small cost makes them ubiquitous. Since the 2000s, most pictures are in a digital format, meaning the information they contain is stored in the form of bits, which are machine-readable and can be processed easily in multiple ways. They can be — and often are — stored on the internet (e.g. on the cloud), whereby they may be privately, semi-privately (i.e. in restricted access) or publicly accessible. Perhaps the most recent development — still in progress — with regards to pictures is the advent of facial recognition software.

Pictures contain — and *are*, arguably — various kinds of data. For example, a digital picture in most cases contains not only the image taken, but also the time and location at which it was taken (metadata). For this reason, in an era where people take a great number of pictures themselves with their own device (i.e. their smartphone), it is valuable to a data controller to have access to one's photo album, because it constitutes a digital agenda describing where the user has been and when. This fact is well-known, and its privacy-related issues have already been addressed at length (see for instance Crandall *et al.* 2010; Loebel 2012; Xu *et al.* 2015).

However, photos do not only yield personal data about the photographer. Indeed, when taking a picture of the Eiffel tower for example, and then making it accessible to a data controller (e.g. posting it on Facebook or uploading it on Google Photos), one provides the latter with material from which personal data about strangers could be extracted. Because someone or face-identification AI could identify the picture's subjects and their occupation at the time and location the picture was taken, one has, in effect, provided the data controller with others' personal data in potency.[35] Thus, like contact lists, pictures are digital material which yields the personal data of a given user, as well as of involved third parties — where this 'involvement' can be unconscious, involuntary or even forced.

European legislation clearly recognises CCTV footage as constituting the personal data of whoever can be recognised therein, and provides data subjects with data protection rights similar to other contexts (EDPS 2009; ICO 2017; Payne 2018).[36] The same thus applies to pictures,[37] except insofar as the (private) photographers will most often not be data controllers themselves — whereas CCTV cameras are mostly operated by data controllers. Despite the

[35] cf. Aristotle's distinction between potency and actuality (Witt 2003).
[36] See also Gras 2004.
[37] Pictures are even considered sensitive data in Dutch law (Parket bij de Hoge Raad 2010).

private nature of the data-capture however,[38] I will argue that N&C should apply to the third-party data subjects, at least when these pictures become processed by data controllers[39] — in which case currently only the user is part of the N&C process, not other people identifiable through the data.

The distinctions made earlier between different kinds of 'identifiable relatedness' apply here too. Pictures will constitute the since personal data of multiple people — importantly, the data of the person whose 'causal agency' made the content available, and of the subjects whom the picture is about ('personal relevance'). That is, if they are recognisable, individuals in the picture also have a relation to it, and the picture thus is (and yields) their personal data. It is thus clear that (the privacy of) the latter should also be protected through N&C.

Now, let us turn to the peculiarity of facial recognition. The kind of privacy-invasive practice discussed with the appearance of photography implies the recognition of a known face from material in a format which can be shared widely and with detrimental effects ('yellow journalism') (Solove *et al.* 2006). While people in a given picture could initially only have been

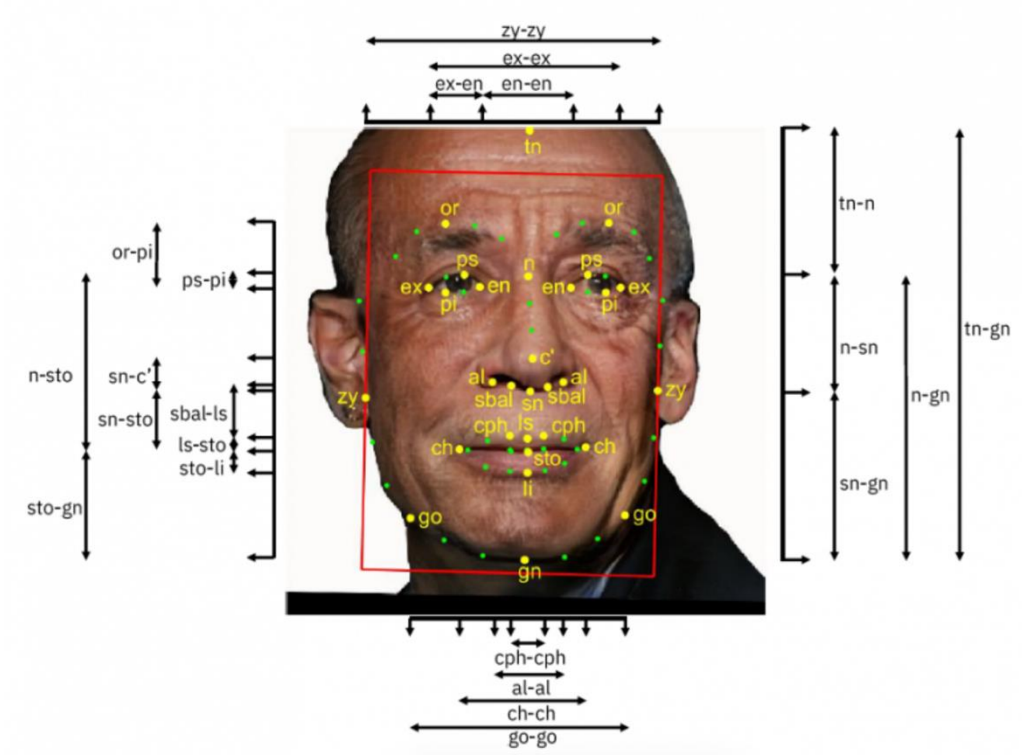

Figure 3: IBM's measurements of a face for the creation of a facial recognition template

[38] In theory, data processing falls outside of the GDPR's scope when carried out by a natural person in the course of a purely personal or household activity. Here however, the natural person asks a data controller to process (e.g. store) the data, thus the GDPR should apply, and *all* the relevant data subjects should have their say.

[39] This is important because, with the data-controller step, private individuals increase the risk that this data, which is not (only) theirs, will be exploited or leaked. This step is moreover important if we want to remain within the scope of the GDPR, which does not apply in the absence of a data controller on which duties of data protection and N&C fall (GDPR paragraph 18).

recognized by a limited number of people, the rise of facial recognition software[40] drastically changes the scope of privacy-invasive photography practices. The increased scope is both in terms of number of people who can be identified by the agent, and in terms of the ease of the process (i.e. of how much material can be processed within the same time).

Computers do not "see" photos and videos in the same way that people do — they just process the data and search for known patterns (see figure 3).[41] Facial recognition AI work in the following way: First, a template of a face is extracted from a photo which is directly linked to a known person — the picture of a mugshot, for example. This template is stored in a database with other templates, which will subsequently be compared to the templates extracted from pictures one wants to identify people in (Article 29 Data Protection Working Party 2012(a):2).[42] If one is not in this database, one's face will usually[43] come out 'unidentified.'

These templates are biometric data and are thus very strictly regulated under the GDPR. This is because

> Biometrics in general are immutable, readily accessible, individuating and can be highly prejudicial. Face recognition, though, takes the risks inherent in other biometrics to a new level because individuals cannot take precautions to prevent the collection of their image. Face recognition allows for covert, remote, and mass capture and identification of images. (Lynch 2017)

This immutability and individuating property of biometrics is the reason why, while pictures in which people figure (and are recognisable) are personal data, the templates extracted from them are *sensitive* personal data.[44]

While face-recognition is infamous for its possibly abusive uses in China or the US for example (Guarino 2016; Privacy International 2017; Burt 2018; Big Brother Watch 2019), its use has been regulated in the EU since 2012 (Article 29 Data Protection Working Party 2012(a-b)). Thus, in the EU the data subject needs to consent[45] before being included in a template

[40] Facial recognition software may take different forms, such as software, algorithms, Artificial Intelligence (AI), etc. I will broadly refer to all these when mentioning about AI.
[41] Source of figure 3: IBM's 'Diversity in Faces Dataset' (2018).
[42] The process includes image acquisition, face detection, normalisation, feature extraction, enrolment and comparison.
[43] Except in cases of misidentification (to which women and people of colour are more subject (Lohr 2018)) partial identification (where multiple matches are presented), or special cases like identical twins.
[44] Moreover, when biometric data may reveal one's ethnic origin for example, it benefits from the highest standards of protection, falling under the category special categories of personal data in the GDPR (art. 9; cf. also paragraph 51, as well as art. 4.14 and 9).
[45] Except for processing on other grounds, such as police's processing in the public interest (but safeguards still apply, and resistance from civil society is strong). Cf. Big Brother Watch 2019.

database; moreover, even if the AI initially processes pictures of everyone, it keeps only those that match the records of those who consented to it (except in cases where consent is not the legitimate ground relied on, e.g. police use), immediately deleting the pictures of faces whose consent it lacks (Article 29 Data Protection Working Party 2012(a):5).

However, more 'regular' uses of pictures are not as well regulated, especially storage. Photos of people may be stored inappropriately, compared to the sensitivity of the data they represent and the impact improper access to them (e.g. because of a hack) could have. This raises three issues. First, they are personal data, irrespective of facial recognition concerns, and should be treated accordingly. Second, some of the subjects concerned (i.e. the people present in a photo) may have consented to be in biometric databases, hence handling their picture inappropriately may seriously impact them (more than those who are not in any database). Third, the threat of facial recognition applications outside of one's control, i.e. without any form of consent or other legal basis, is real and concerns everyone (Guarino 2016; Chen 2019).[46]

Because pictures are so common and widespread, and because most often no strict rules apply to them when they are in the hands of private individuals, they are taken without much consideration for the people they capture, and are handled very nonchalantly regarding appropriate storage and protection. People regularly store them on the free online storage services (such as Google Photos, iCloud, Facebook Albums), thereby agreeing that the storage provider acquires certain rights[47] with regards to such content. Pictures are thus stored online or on personal hard drives, which may be compromised: hard drives may be stolen, photos shared publicly on Facebook may be appropriated by others,[48] and cloud storage may leak.[49]

The point is, photos of people are personal data, and the people whose personal data they represent can increasingly be affected by any possible mishandling of their personal data in this form — with data-mining techniques and facial recognition AI for example. While these photos are multiple people's personal data, N&C currently only covers a single person's rights

[46] IBM for instance used one million Flickr pictures (uploaded by private individuals under a Creative Commons licenses) to fuel the development of its facial recognition AI (Burt 2019; Liao 2019). Although none of the relevant data subjects may have been harmed or their privacy violated, many expressed concerns about their indirect participation in the technology, which is known to be applied in controversial uses—such as IBM's helping Philippines' president deathly repressive practices (Joseph 2019).
Another example is Google Glass, the 'smart glasses' which in 2013 raised concerns in the public regarding covert filming and facial recognition uses (Livingstone 2013; Henn 2013).

[47] Data controllers thereby often gain, not ownership of the content, but many rights and leeway regarding the use they can make of it.

[48] See for example http://stopstealingphotos.com/ (accessed 02/06/2019), a website for photographers whose work is being stolen and appropriated by others through social networks.

[49] Cf. iCloud's 2014 protocol flaw which allowed people to get access to others' accounts and steal their photos — dubbed "the fappening" (or "celebGate") because of the numerous celebrities' nudes (Lewis 2014).

to privacy-protection. The behaviour of this user is designed to be convenient (e.g. storing or sharing pictures online), but it often generates a foreseeable privacy externality, a price paid by others.

I will now examine how we can ensure that due care is taken to protect others' privacy, proportionally[50] to what is at stake.

[50] Cf. GDPR's article 14, which mentions efforts that are "disproportionate" to the data controller's duty to apply N&C in certain cases, as well as the corollary required "appropriate measures to protect the data subject's rights and freedoms and legitimate interests" when this is the case.

4. Addressing Interdependent Privacy

In this fourth chapter, I first put the N&C model to the test to see whether it can meaningfully be applied to the cases discussed above. I proceed from the assumption that it would make sense to apply the N&C model here, insofar as these cases of privacy externalities represent clear interferences with the ideal of privacy self-management. Because it quickly becomes evident that a straightforward application of the N&C model alone would be very burdensome as well as disproportionate, I then explore three hypothetical strategies for the development of privacy-protection.

4.1. Applying the N&C Framework

The N&C model prescribes that data controllers notify all data subjects and request their consent (for specific, pre-determined purposes) before[51] starting to process their data. However, this should only concern those subjects who already want to benefit from the services offered by the data controller, such as someone initiating a request to use its services (the user) or individuals already using these services (existing users). Allowing the controller to send its notice to just *any* person whose data it got its hands on (e.g. through data leaks) is not allowed; it would amount to (illegal) advertising,[52] and would burden the myriad of people whose data is leaked or disclosed through no fault of their own to reject the service offer. It would moreover create perverse incentives for data controllers to offer their users precisely such services that disclose other people's data.

Therefore, upon reception of material containing multiple people's personal data, the data controller should first compare it with its own database to see whether some of the third-party subjects affected have not already consented to their data being processed by the data controller. When this is the case, following N&C we still expect the controller to approach these data subjects to inform them that someone disclosed their personal data[53] — especially if the controller thereby acquired additional data, the processing of which the subject may oppose to.

Regarding the processing of personal data of subjects whose consent it lacks, the data controller needs to be particularly cautious. Indeed, even the mere comparison of this data with

[51] Or, at the latest, right at the beginning of the processing if N&C cannot be done otherwise.
[52] A cheap and effective way to reach out to new customers.
[53] Facebook arguably made a step in this direction with its 'Custom Audiences' transparency feature, which provides the user with information regarding the origin of an ad and the company that uploaded her contact information for targeting. This only addresses contact information uploaded by advertisers however, not to one's personal data disclosed by others in general. Cf. Constine 2019.

the controller's own 'database of consenting data subjects' amounts to unrequested (thus illegitimate) processing according to N&C. Precedence with facial recognition regulation could apply here,[54] as such processing could be tolerated insofar as the data is used only for this purpose of ascertaining the existence of consent, and immediately deleted if no overlap is found with the data controller's database of consenting users.

Therefore, if we want to be consistent with our current model of privacy-protection and apply N&C to the cases of interdependent privacy, then there should be a mechanism whereby the data controller, upon reception of material by an informed user, 1) compares it with existing databases to determine overlap, 2) notifies existing users of the disclosure of their data, and request their consent for its processing if new data is involved. It would then either 3a) extract from the material the data it may process, and delete the material itself (e.g. the contact list, which contains un-consenting third-party subjects' data) or 3b) keep the content but delete or fully anonymise the data of subjects whose consent it lacks (e.g. render certain faces unrecognisable).

Two things should be noted already. Firstly, the choice between 3a) or 3b) will depend on the nature of the content and the use that can be made of it by the controller. For example, if for whatever reason it is valuable to the controller to keep a record of each user's contact list in such a form, 3b) will be preferred; but if the sole purpose of the contact list is to access the users' contacts and determine at once who the user can be connected to, the controller should favour 3a), because it then does not need the list itself anymore.[55] Secondly, in the case of uploaded contact lists, note that due to the necessarily interpersonal nature of the material[56] the data controller already knows from the onset that the content it receives contains the data of subjects other than the user — data which is already attributed clearly to distinct data subjects. The example of a contact lists upload is helpful because it is very clear in this regard, unlike other materials for which it may be less obvious or expectable that multiple people's data are included. If we want to apply the aforementioned mechanism to other contexts and materials than cases of necessarily interpersonal data, we may want to add a preliminary step that should be performed: 0) analysing the content received to a) determine its nature and whether it

[54] In facial recognition software, newly acquired templates are allowed to be compared with those existing in the database, only the matching of which triggers the AI to react. When no matches are found, the new templates are deleted without further ado (Article 29 Data Protection Working Party 2012(a):5).
[55] In the latter case, not opting for 3a) could perhaps be said to be in contradiction with the controller's duty to only keep data it needs for the offering of its services (data minimization principle, GDPR Art. 5(c)).
[56] Contact lists provide relational data which allows the service provider to connect the user with people she already knows. In so doing, the controller necessarily knows that subject 1 (the user) knows subject 2, and can reasonably infer that subject 2 also knows subject 1. Contact lists moreover directly hand the personal data of the user's contacts over to the data controller.

includes personal data of other data subjects than the user, as well as to b) assign the different data points to their relevant subject.

This mechanism seems technically feasible, but particularly burdensome if carried out by a human — especially this preliminary step, which would for instance require reviewing the thousands of hour of video uploaded on YouTube every day. It implies that for the sake of N&C, data controllers always ought to analyse *any* content they receive, because it might always contain the personal data of people whose data it cannot legally process. Indeed, even merely storing the content without analysing it would amount to 'processing' the personal data of potentially un-consenting subjects (GDPR Art. 4.2.). Given the magnitude of the task, this mechanism obviously is no task for humans: algorithms will have to do it.

If carried out by algorithms, this would be similar to the existing means used to detect violent, hateful and sexually offensive content, as well as spam, copyright-infringements and so-called 'fake news.' None of the algorithms performing these reviewing tasks is fully accurate, and the same is foreseeable of the task proposed for the sake of N&C. Instances of privacy externalities might not be detected, or conversely many might be falsely detected. The mechanism proposed moreover seems very similar to the 'upload filter' recently discussed in the EU through the notorious 'Article 13,' which will in effect require platforms to algorithmically analyse all uploaded material to prevent copyright infringements, for which they would be liable in their quality of host of the material (Ferrer 2019).

Besides not being wholly accurate, the mechanism I proposed could thus also share the flaws of the copyright-detection algorithm, the most notorious of which is the plausibly negative impact on freedom of speech it will have. This is because platforms will likely be overly strict and censor material that is fully legitimate, in the fear of missing instances of copyright infringement and incurring the resulting fines (EDRI 2017, 2018, 2019). Similarly, the logic behind the mechanism I offered above could protect privacy at the cost of impeding on other freedoms and societal benefits.

In conclusion, it is one matter that consistently addressing privacy externalities would already greatly impact our lives: it would affect the things we do online (e.g. sharing pictures), the digital services that serve us offline (e.g. smart homes), and the activities that affect the privacy of multiple persons (e.g. genetic analysis). It is another that, if we use N&C alone to address privacy externalities, the only feasible way to address the colossal amount of material would be to adopt algorithmic mechanisms, which could foreseeably restrict us even more than is necessary — even though what is necessary is already a lot. Our pictures could be forcibly blurred before they would be shared on social media (see cover), blogging and other forms of

expression would be monitored, genetic analysis would be prohibited, etc. Ultimately, we could miss on crucial societal benefits, as innovation would be stifled, free speech[57] restricted and health research slowed down. Applying N&C to privacy externalities has stringent consequences, and relying on N&C unaided is even worse, as it appears to have disproportionate consequences. The precautionary principle, which has become part of the European spirit,[58] states that when a given serious harm is foreseeable, the magnitude of which is unknown but plausibly significant, measures should be taken to prevent it. While I cannot quantify the risk and magnitude of the harm aforementioned, these appear significant enough to warrant applying alternatives or complements to N&C — or at least significant enough to warrant even just exploring these alternatives and complements.

I will now explore three strategies, the first and third of which in effect complement N&C, while the second stands as an alternative to it. They are meant to ease the tension between our value of privacy-as-control (implemented through N&C) and our practices (which sometimes generate privacy externalities). Perhaps can the burden and impact of the mechanism I proposed be rendered less disproportionate through the help of some specific technological fixes — fixes which would allow us maintain our externality-generating practices while upholding our value of privacy-as-control (strategy 1). Perhaps we need to relinquish our concept of privacy as control, however, because we could thereby avoid the disproportionate impact and burden that follow from applying N&C consistently (and exclusively) (strategy 2). Finally, since the systematic application of N&C is too burdensome according to our current standards and practices, perhaps it is the latter that needs to change, because the value of freedom, instantiated in the concept of privacy-as-control, is more valuable (strategy 3).

4.2. Strategy 1: Technological Fixes

The problems we face regarding information privacy, which is a concept brought about by advances in technology in the last 50 years, could perhaps be solved by *more* and *better* technology. Technological fixes could be designed to lighten the burden of privacy management discussed in chapter 3 (so that privacy self-management becomes truly possible) and to minimise the impact of consistently applying N&C (so that cases of privacy externalities are duly addressed). This would, in effect, not to change our expectations about privacy nor

[57] 'Speech' here refers to any form of expression, including sharing pictures.
[58] Cf. Recuerda 2006:283; Von Schomberg 2012.

our attachment to individual control. This is hence an ideal scenario, because it is one in which little change in expectations, standards and practices has to take place.

I will now introduce two such kinds of technological solutions, one which has been offered in the literature already, the other as my own. Both will be technologically feasible, but will face serious shortcomings and issues (distinct from the shortcomings of the mechanism described above, which they are meant to complement). I will hence cast doubt upon whether they are likely to be (effectively) implemented, and in light of this uncertainty, I will consider two alternative strategies which we should choose from until these technological fixes happen — if they ever do.

The first kind of solution examined is the 'outsourcing' of the labour of standard privacy self-management. This can be done to various degrees — from easing the cognitive load to automating decisions — all of which are applied through the establishment of some kind of interface where one's privacy can be managed. A good example of this is Pascalev's (2017) Privacy Exchange Authorities ('PEA').

Pascalev's concept of PEA is that of a centralised privacy self-management platform. This PEA allows the data subject to choose once (and to periodically update) what processing of her data she will allow. This is based on information and guidance provided by the PEA, as well as carefully-considered trade-offs. The decision for each kind of personal data and each different use of it is then stored by the PEA, and a report of these decisions is issued to the data controller every time the data subject wants to use its services.

Thereby, the contract between data controller and user would become more like a bargain — each party comes with its own terms, and if no middle ground can be found (e.g. the user's preferences are too strict and the controller's preferences too broad), no 'contract' is made, and service is denied. The advantage of this system is that the user can make an informed choice once, and have it more or less consistently applied every time after.

Proposals similar to Pascalev's have been discussed elsewhere (de Montjoye *et al.* 2014; Poikola *et al.* 2015; Obar 2015:13; Chaudhry *et al.* 2015; Obar 2015; Lehtiniemi and Kortesniemi 2017).[59] For the sake of convenience, I address all these as forms of PEA. Fundamentally, every PEA aims to make privacy self-management easier regarding (some specific) shortcomings of the model — to the extent that it remains 'self'-management when external parties are involved and when choices become automated. They also centralise self-

[59] Implementations of such PEAs can be found in the following applications: Hub of All Things 2017; Meeco 2017; Digi.me 2017; Cozy Cloud 2017.

management, which allows data subjects to have an overview of how many invasions of their privacy they have allowed. This sort of solution is extremely attractive and, in my experience, it is also the one people think about first when confronted to the burden of privacy management. Indeed, as they automate and ease the privacy-protection processes people go through, PEAs seem to only require the creation of a platform and little (if any) change in individuals' practices. They are both simple and convenient.

However, PEAs face their own issues, and create problems that did not exist beforehand. I will here briefly criticise them in six points.

1. It is difficult, if not impossible, to standardise and codify the privacy-related choices in a meaningful way; these choices need to be streamlined and (often) binary so as not to create a mess of individual preferences for data controllers (i.e. to maintain the societal value of data), and the PEA-like solution will not capture many nuances between different choices and behaviours otherwise available.[60]
2. Insofar as we intend to retain the concept of self-management as central to privacy-protection, a PEA will not necessarily increase the proportion of privacy self-managers. Some individuals will not reflect at all about the different possibilities, and will simply choose one option (e.g. for convenience the one allowing for the widest range of services, and hence the most intrusive practices) then let the PEA do its work. This would make us question to what extent this system retains its core value of control, as well as to what extent the user provides informed consent in contracts with data controllers.
3. These systems cannot be offered for free as they require important investments for their creation, whereas the existence of the "privacy paradox" shows that most people are not willing to pay (or make additional efforts) for increased privacy (Kokolakis 2017; Barth and de Jong 2017). If they are not subsidised by governments or by big tech (something which creates distinct issues, see point 4 below), there will thus likely be no widely-accessible PEA. Moreover, a majority of services of all sorts (i.e. both websites, smart watches, connected cars, genetic research labs, etc.) will have to adopt this model, for it to be deemed attractive by data subjects and to be deemed worth the costs of adoption by data controllers. This

[60] This is acknowledged by Pascalev (2017:44), who mentions that this will require the cooperation of all major stakeholders: legal scholars, ethicists, technology practitioners, privacy advocacy groups, government agencies and big data companies.

did not happen with the abandoned web protocol P3P — to which Pascalev's PEA proposal is very similar.[61]

4. Relying on these systems will create new concerns about power, regarding for instance who designs the algorithms automating decision-making, who designs the choice architecture, which country hosts the database, etc. Power can be abused — especially when centralised — and backdoors created, and such fears will be directed to whichever company or state builds the system. However, even without abuses these PEAs will constitute troves of information and centres of influence likely to be the target of major cyberattacks (Lehtiniemi and Kortesniemi 2017:10).
5. Finally, while some shortcomings of the privacy self-management model will likely be solved, and others mitigated by these solutions, others will remain — particularly negligent behaviour and privacy externalities. Fundamentally, the framework of PEAs remains extremely individualistic: sovereignty is still implicitly granted to the data subject regarding the use of the data in her possession, overlooking the fact that it often implicates third-party subjects.

This last criticism is crucial, because even if we solve the other shortcomings of the N&C model through PEAs, we cannot meaningfully be said to have privacy self-management if others continue to (inadvertently) interfere with it. I think that a PEA can be specifically designed to also address privacy externalities, however. Let us now consider such a system, with the case of pictures, where we will see some of the other criticisms return.

To rephrase what has been said before, pictures can constitute the personal data of multiple people at the same time, insofar as these people can be recognised therein (by a machine or a human). While the action of taking a picture for private purposes is out of the scope of the current legal requirements for N&C, allowing commercial entities to process them is not (where 'processing' means any action performed on the picture, from storage to analysis).

Naturally, the user of a storage service like Google Photos could provide at the time of upload the contact details of each person in the picture in order for the data controller to conduct N&C. Secure mechanisms could even be designed so that this contact information is only used for the goals of N&C and deleted afterwards. However, this system would not be sufficient to

[61] P3P stands for the Platform for Privacy Preferences Project, a protocol (or machine-readable language) that expresses one's privacy preferences; used by both website and user, it would have automatically compared the terms on each sides and determined their (in)compatibility, only prompting a message and thus privacy *self*-management when further action is needed. Besides poor adoption rates — privacy was not as prominent in 2002 as it is in 2019 — P3P also suffered from criticism over its security. Cranor 2002; Stufflebeam *et al.* 2004.

ensure suitable protection to everybody, because — among other things[62] — the user is likely not to know everyone included in her picture, especially when taking a picture of a tourist attraction abroad, for example.

The adequate solution for solving this challenge is thus perhaps the creation of a worldwide biometric database. One would willingly add oneself to this database, which would be solely used for the purpose of notifying and securing the consent of those whose personal data are put online in the form of, say, photos. Any photo put online would be processed, and the templates sent to the database for comparison (and deleted afterwards). Consent requests would be sent to the people recognised by the facial-recognition AI, with information about the processing (e.g. storage or data mining) and its goals (e.g. private or commercial), as well as about the nature of the data controller (e.g. Google) and perhaps that of the user (who uploaded the picture) too.

Assuming perfect accuracy of facial recognition (i.e. no false positives) and an adoption of this method that is widespread enough, the event of a detected face not finding its match in the database would effectively mean that the relevant data subject does not consent to the processing. When consent would be passively lacking or actively refused, the relevant personal data would be anonymised. Anonymisation could be achieved by blurring the face of the un-consenting data subject. However, this would not only ruin many pictures (in the sense of altering their artistic and social value), it would also not always fully anonymise the un-consenting data subjects. Indeed, other clues than faces (especially clothes) together with contextual indicators (especially timestamps) may be sufficient to precisely determine the identity of data subjects, meaning that adopting the 'blurring' strategy would require the blurring of potentially important portions of pictures — such as whole bodies (Song and Leung 2006; Gallagher and Tsuhan 2008). To avoid identification through complex means, as well as a blow to the field of personal and professional photography, AI recognition could be misled by editing the picture with an artificially-generated face. Nowadays, an algorithm known as a GAN can render hyper-realistic portraits of completely fake people;[63] these could be used to match the posture and attributes of third-party subjects, and replace their face inconspicuously — thereby not only anonymising them, but also misleading an observer (machine or human), for whom or which the face is the main identifier.

[62] Another issue is that the user will not necessarily want to have the approval of everyone in her picture before uploading or publishing it online, especially if it's only for her personal use. Fundamentally, this system would still let third-party subjects at the mercy of the user's willingness to apply third-party subjects' desires, hence maintaining the power imbalance in privacy management.

[63] See the generative adversarial network (GAN) generating faces at https://thispersondoesnotexist.com/ (accessed 02/06/2019).

Relying on such a technological system would have certain benefits. Currently, when a user enters an agreement with a data controller, the concerned third-party subjects are not aware that their data is processed, and therefore have no way to express consent or dissent.[64] Moreover, *if* they encounter pictures of which they are a part — such as on someone's Facebook wall — they lack effective control (and hence, privacy self-management) because they usually can at most "give feedback or report" the content. Most personal data posted about us by others may pass unnoticed anyway. The method here proposed would thus have the advantage that, unlike the current way of handling data, the third-party subjects would either be represented in the N&C process (by being sent a consent request), and their data would be automatically anonymised were they not to agree, or if they were not able to be contacted in the first place.

While it has its benefits, this method also raises its own issues. People could opt-in for the sake of always being alerted of their picture being processed or posted online, but the cost would be very privacy-invasive itself. A hack, leak or backdoor in such a gigantic biometric database would be immensely more problematic than the leak of Facebook's passwords, for example, because unlike passwords, biometrics cannot be changed, and can be used for identity theft or fraud for instance.[65] This database could also be abused or repurposed by the authorities of the jurisdiction in which it is based — a foreseeable instance of 'control creep' — which will incentivise different countries to establish their own database, and will weaken the overall strategy (which should be based on a single database if people are to be protected worldwide) (Innes 2001; Kitchin 2014a).

Moreover, current face-recognition AI is still imperfect, and especially mis-identifies women and people of colour (Lohr 2018; Buolamwini and Gebru 2018). This means that exclusively relying on this database could hypothetically lead to people being mis-identified and allowed to manage the privacy of others. Many projects of facial recognition or biometric databases are being challenged for all these reasons, by civil society or even the employees and shareholders of the company developing the tool (Electronic Frontier Foundation 2012; Dent 2019; Kelley 2019). What we are advertised is a technology that might lead us to some sort of utopia (where being a fully-autonomous data subject, managing one's own privacy, and addressing privacy externalities is effortless), but given the current development of such

[64] These issues of awareness and power are a tough issue which regularly limit the usefulness of other proposed technological fixes, such as Ang *et al.*'s (2017).
[65] Biometric can be changed to a certain extent, but not in the same way. For example, face surgery could do the trick here, but is very expensive. For those who cannot afford it, just one breach of the biometric database is enough to ruin things for life.

technologies and in light of other consideration such as concerns about abuses, what we are sold is a technology that risks very much leading us to some sort of dystopia instead.

These issues have the consequence that it is not clear whether we should wish for such a system in order to better protect our privacy, and whether it would be a proportionate solution; maybe the price to pay (the danger of a worldwide biometric leak) is too high already, compared to the intended benefits (actual privacy self-management despite other people's negligent behaviour). Moreover, together with some of the issues raised for the above tech fix (especially 3-4-5), the issues faced by such an infrastructure mean that, while this solution is technically feasible, it depends on various demanding factors (especially trust in the technology and a very widespread adoption of it) which will need to be present together for it to work.

Finally, as this solution is specifically tailored to pictures, it means that similar tools should be designed for the other cases, which may be less straightforward to address. If the goal is a holistic form of privacy self-management, many instances of negligent and inadvertent meddling with others' privacy would need to be fixed separately, each of which would encounter its own obstacles. This is because what should perhaps be framed as a complex (or wicked)[66] societal problem of 'negligent behaviour,' 'unforeseen interdependence and externalities' and 'overlooked responsibility' (cf. strategy 3 below and its duty of care) is instead reduced to an issue which can be solved by applying "quick technological fixes," which address the symptoms but not the root cause (Weinberg 1966).

Thus, we should perhaps revise our expectations while we wait for these tech fixes to happen — if they ever do — not only because the technology does not look like it is ready, but also because it might also bring with it problems more serious than those which it was there to solve in the first place. That is, it solves the problem of burdensomeness of N&C, but not that of disproportionality. The reason we considered strategy 1 is because it would be extremely convenient to continue living as we do while upholding a model of privacy-protection that matches our valued individual freedom. In this case, our way of life generates privacy externalities, and our model of privacy-protection is based on the value of individual freedom, in the sense of privacy-as-control applied through N&C. We might be better off revising either of these two however, because I showed that applying N&C to privacy externalities would have disproportionate side-effects, even with technological complements. Let us now examine these two strategies in turn.

[66] Cf. Morozov 2011:281.

4.3. Strategy 2: Renouncing Privacy-as-Control

Perhaps the easiest thing to do, given the difficulty in addressing the challenges to privacy self-management with tech fixes, is to reconsider what privacy fundamentally is, or, alternatively, how we should protect it. To reiterate, (information) privacy has, in this thesis, been understood as the relative control over the ways and the extent to which one selectively discloses (information about) oneself to others. Considering the challenges to the realisation of this concept we have encountered and detailed in the preceding sections, it might be preferable to try living without (this kind of) privacy.

This narrative fits two main strategies, only the latter of which I will discuss at length. One strategy would be to abandon privacy altogether, in the spirit of those who proclaim the "end" or "death" of privacy (Whitaker 2000; Enserink and Chin 2015, Solove 2008:5). This cynical, deflationary strategy would argue along the lines that, while we may have had meaningful means of privacy for about fifty years, the factors discussed in chapter 2 (especially the growing power of large corpora) have made privacy in the last decade impossible, and that it is not a great loss. On this view, most of us cannot have privacy anyway, and it is pointless in the digital to pursue era the ridiculous ideal that one can have much privacy. While I mention this strategy because it may be socially relevant (given the lack of interest some are thought to show regarding privacy (Barth and de Jong 2017; Kokolakis 2017)), I will hold to the assumption made earlier (that privacy matters). Therefore, the revisionary strategy on which I will focus instead addresses the very definition of privacy and re-thinks the foundations of the concept.

The concept to be amended in this section is that of 'privacy-as-control.' What is at stake here is the individualistically-framed value of (freedom, conceived as) control which, as discussed in chapter 2, has been a core component of the concept of privacy for a long time in the West. There are multiple ways to redefine privacy, thus for the sake of illustrating the kind of changes that this strategy would entail, let us make a thought-experiment about the possible world 'Enor,'[67] in which the degree of control people have over their privacy is much smaller than what it currently is for our world. It is not inexistent — control simply is not the core of the concept of privacy.

In Enor, privacy is centred on intimacy. It is not framed as 'control over the extent to which one discloses oneself selectively to the world,' but rather as 'being respected as a human person in the aspects of one's life that should remain private, i.e. free from the gaze of the

[67] The name Enor has been chosen for its etymology: it is the ancient form in French of the word honour, which is tightly linked to the way privacy will be conceptualised here.

other.' These aspects could include, for instance, one's financial and health status as well as photos of oneself in the confines of one's home, but exclude information about what one has said online or information collected within the public sphere. The aspects that are constitutive (or not) of individuals' privacy are decided by the state and not by the individual, although they are partly based on pre-existing social conventions — which are adapted to the new situations enabled by technology — and although the government is democratically elected. The universe of Enor is not very different from ours in terms of practices: people capture and share pictures, their personal data and that of others, without thinking about their and others' privacy. What is different is the values on which these practices are founded.

In Enor, people are not notified about the existence of data processing, nor are they required to consent to it. They do not have to care about ensuring their own privacy themselves, because the matter of privacy is dealt with by experts, on behalf of the whole population. These experts are specialists in the cultural norms of intimacy, and they are the ones to allow or forbid certain uses of personal data.

In Enor, citizens merely expect that nothing 'bad' will happen, because the adequate processing of personal data is defined and enforced by the state's experts. Data controllers in Enor, as well as their systems and practices, are regulated and audited. This happens just as the state in our own, actual world regulates the construction and provision of goods and services, as well as it ensures compliance to the rules through a system of inspections, diplomas and seals of approvals. Data controllers in Enor do not rely on the N&C framework, they simply go ahead and process the data available to them within the scope of their services. For instance, while a 'smart' supermarket in Enor cannot communicate the content of its customers' (healthy or unhealthy) purchases to their health insurance, it is allowed to internally process this data to train its personal assistant AI, because the latter is considered a 'fair and legitimate use' of the data by the state's expert.

While individuals are not expected to take care of their privacy, there are rare instances where they suddenly become aware of a lack of privacy — that is, of intimacy — because a data practice allowed by the state's experts actually makes them feel uncomfortable. In such cases, if enough people flag the data practice as uncomfortable or inadequate, a case is brought to the experts, who may reconsider their original decision. There thus exists a certain degree of control, but it is far from the degree we are expected to exercise in our actual world.

The situation in Enor is not far-fetched. It is certainly more paternalistic than the state of privacy-protection in the EU in 2019, but it is a possible world in which we could very well

have ended up. It is in the direction which several critics of the N&C model of privacy-protection advocate for, such as Burt and Geer's (2017) opinion on "the end of privacy,"[68] who say that "the future of our privacy lies in how our data is *used*, rather than how or when our data may be *gathered*. Excepting those who opt out of the digital world altogether, controls on data gathering is a lost cause."

The situation in Enor has clear advantages and disadvantages compared to our current, actual situation. It is less burdensome overall, privacy-wise: individuals do not have to care about ensuring the data controllers whose services they rely on are not unfairly exploiting them — just like we do not have to care about ensuring the food we buy in a restaurant will not make us sick, because there are health and safety standards which are regularly checked and enforced. Moreover, in Enor privacy regulation addresses the usage of data, in addition to our current situation where what is regulated is mostly its collection.

The reason this thought experiment was rolled-out was to explore graphically what a model of privacy-protection alternative to self-management could look like, because N&C appears to be excessively burdensome and to have disproportionate effects. However, strategy 2 as described in this thought-experiment might face stringent issues. The trade-offs from relinquishing our ideal of privacy-as-control should be well-examined, and safeguards installed. Part of the crucial benefits of having privacy(-as-control) are that it may be a precondition for critical thinking, independence and, ultimately, democracy (Allen 2011:21-2; Cohen 2012:144). In a world where the state holds centralised authority over what is appropriate and what is not, the shadow of authoritarianism may loom close. Naturally, this risk is thwarted by the safeguards put in place, but even in a democratic state this centralised power can harm — for instance if the majority allows uses of data that are specifically harmful to minorities.

Strategy 2, just like strategy 1, seduces us with an appealing scenario of a 'convenient lifestyle.' Here too, the utopia could turn into a dystopia. An important degree of trust would need to be at play in these strategies for us to adopt them, as well as carefully-defined safeguards. Again, however, let me highlight that these conditions can be met, and that Enor is only a non-exhaustive exploration of a specific alternative to privacy-as-control. As less paternalistic alternatives exist, we need not go as far in the direction of paternalistic interference in our private lives as Enor goes.

[68] For a similar position, see also Baeles and Muris 2008; Tang *et al.* 2008; Nissenbaum 2011; Center for Democracy and Technology 2018. Cf. Morozov 2013:299.

The Enor thought-experiment was constructed to help examine what relinquishing our commitment to (the concept of) freedom-as-control could look like, while maintaining our negligent way of life. Let us now consider the opposite strategy, wherein we uphold the concept and discontinue some of our practices.

4.4. Strategy 3: Biting the Bullet

Privacy-as-control is burdensome, and the mechanisms and adaptations proposed in strategies 1 and 2 are improbable or risky. The strategy to be examined here is one in which we decide to bite the bullet on demandingness: on this view, we should not question our attachment to the value of privacy and of control, just because consistently applying these values would mean burdensome changes to our way of life. Yes, the full downstream implications of adopting the N&C model are burdensome, but it is our behaviour we should modify, not our ideals and core concepts. Ethics places demands, and we will not flinch. Moreover, the disproportionate effects of consistently applying N&C to cases of privacy externalities can be avoided by making sure these are not 'externalities' in the first place, i.e. by internalising the costs.

The value being traded-off in this strategy is thus that of convenience; it has hitherto manifested itself in the form of burdensomeness. It is relevant in the equation because it is omnipresent, from the efforts we are already expected to make to control our privacy, to the efforts we would be expected to make to protect that of others, to the changes that alternatives to N&C require. It would be difficult to set the hypothetical threshold from which a given goal or value could be said to be too cumbersome to implement. Instead, I will again make a thought-experiment (the possible world 'Unith'[69]), this time where people are serious about privacy — theirs, as well as that of others — and about their commitment to individual control. Again, the purpose of relying on a thought-experiment is to give a sense, experimentally, of what following this strategy could look like.

In Unith, people allocate a certain amount of time per week to the examination of privacy policies, in order to carefully select those services which lead to reasonable trade-offs. People are educated at school about the intricacies of privacy self-management, and are then expected to manage their privacy themselves — just as they are expected, from 18 years of age, to fully read and understand contracts before signing them, for example. Individuals are thus well-aware of the trade-offs they agree to, when they do. While not everyone cares as much

[69] The name Unith refers to the possible world's unity (concordance) between their values and their acts.

about their own privacy, everyone does care about not infringing on the privacy of others — this is what is socially and legally expected.

In Unith, people do take pictures — even of the public sphere — but when they capture the face of others, they often ask their consent beforehand. When that is impossible (in the case of photographing moving crowds, for instance) or when they deem it too burdensome, they still blur parts of the picture to anonymise un-consenting data subjects before sharing the picture online or submitting it to the processing of a data controller. They typically genuinely care about not violating the privacy of others, but when, for some reason, that is not respected, they are reminded by others. There is thus in Unith a societal expectation of respect of each other's privacy, a sort of duty of care for others that is so important that it extends to their privacy. Therefore, what was previously an externality is internalised, in the form of additional efforts and of lower quality of practices — such as taking the time to blur certain parts of one's pictures.

People can take out insurance for their privacy and for the harms resulting from damage done to it,[70] and insofar as one handles other individuals' personal data (and thus potentially damages their privacy), one incurs the legal responsibility of taking due care. Moreover, while the model of N&C still applies in the form of a contract, it is far more inclusive than the model in our actual universe, as contracts always take third-party subjects in consideration besides the user. Regarding the disclosures of third-party subjects' data that are unavoidable (such as talking about someone's private affairs in an online conversation) either it happens through the services of a data controller that is known to have committed to not exploiting such data (one which, for instance, has implemented P2P encryption in online communications), or the people disclosing such data do so knowing that they incur the responsibility for what happens to the data. That is, they are liable for the damage their acts may bring upon the privacy of others.

In Unith, great care is taken with material which affects the privacy of multiple persons and for which it is impossible to notify and gather the consent of every affected data subject. Most often, this material is simply altered to anonymise those who did not or cannot provide consent; sometimes however, this material is simply destroyed. This is the case of commercial genetic analysis results, after they have been produced and directly used on the 'user' — and only for purposes that affect her only. This is because genetic records affect the privacy of the whole chain of dead, living and future descendants, which cannot (or not without difficulty) be notified or give consent. Indeed, as already noticed by Manson and O'Neill (2008:119), "all

[70] For instance, MacCarthy (2011:32) identifies five possible harms resulting from indirect disclosure of information: invidious discrimination, group injury, inefficient product variety, restricted access to products and services, and price discrimination).

DNA information is potentially linkable, none of it can be irreversibly anonymised," and requiring the consent of all relevant actors "would also put an end to virtually all genetic studies using lawfully held samples and data." In cases of doubt, the precautionary principle therefore rules in Unith, and the material is either destroyed before it can harm, or not created in the first place.

Just like with the GDPR however, a certain balance between privacy and the societal benefits from processing of personal data is sought in Unith: while commercial processing of data is tightly restricted by the duty of care mentioned above, the private, public, academic, journalistic and artistic fields are less restricted by it. They remain very concerned about privacy however, and the data of their subjects they examine, cover or address is extremely well-protected.[71] This allows a certain degree of research and press coverage, and prevents a too slow rate of scientific development (compared to the situation where they would have a private market for research), these fields are well-funded by the state. This allows people to benefit from advancement in fields where the privacy of many is impacted (such as genetic research) without risking the dangers faced in the current, actual world where corpora are increasingly getting their hands on all sorts of data.

Concretely, a transition to Unith would be very burdensome. It would mean adopting new standards, enforcing legal and social expectations, as well as abandoning certain practices that necessarily impact the privacy of others. It would increase the complexity of privacy management, and require considerable efforts from individuals. It would be the end of carefree photography, and it could still have important consequences in terms of slower societal development. It would also somewhat impact freedom of expression, because of the increased complexity of carrying out certain practices, and because criticising others online would be less easy — we would be liable to a greater extent for defamation, for example.[72] This casts doubts about the proportionality of this strategy compared to the issue at hand. This transition to Unith appears very burdensome[73] only *in light of* our current expectations and practices however, i.e. in light of the relative convenience of our current way of life. Perhaps it is the latter which is in the wrong: perhaps we should have been more considerate of others' privacy from the start, perhaps deeply caring about privacy should be the norm.

Unith was sketched in order to explore what we should amend in our way of life to ease the tension between our commitment to privacy-as-control and our externality-generating

[71] Cf. GDPR paragraph 153, and art. 85.
[72] Cf. GDPR art. 85.
[73] This is a demandingness argument (cf. for instance Hooker 2009). Given we are here trading privacy (a fundamental right) off, this argument only has normative ground based on the proportionality and plausibility of the solution and its alternatives.

practices. The situation in Unith appears more burdensome than in Enor (but it also appears less demanding of a change, for what matters). While in Unith privacy-as-control is still central, the problem of privacy externalities is treated at the source (with a duty of care), preventing them instead of treating them through N&C, thereby precautionarily avoiding the disproportionate side-effects of consistently having to apply N&C.

5. Blending the Strategies: A Suggestion

In summary, we have examined three strategies describing possible directions to follow alternatively or complementarily to N&C. They were framed as upholding either or both of the values 'freedom' (insofar as instantiated in the concept of privacy-as-control) and 'convenience.' Crucially, they were intended to address privacy externalities in a way more proportionate than what relying on N&C alone affords. However, ultimately, while strategy 3 fared well on this aspect by confronting the source of the problem (the duty of care we currently do not perform), strategy 1 and 2 were more successful in tackling the burdensomeness problem.

These strategies have their benefits and drawbacks; the duty of care from strategy 3, for instance, has the advantage that it addresses particularly well the problem of privacy externalities, but at the expense of increasing the burden of privacy management for everyone. It is possible to view the directions these strategies take as compatible, however — especially if they are only complementary to N&C (strategy 1 and 3), not alternatives to it (2). Given that none of them seems able to solve privacy externalities on their own without potentially disproportionate side-effects, observing the precautionary principle might require from us that we mix them and take the best of each world — and this is what I suggest here.

By framing from the start the issue of privacy externalities as an morally-loaded social issue rather than simply as a technical issue of excessively-wide sharing of data, I have also framed what, accordingly, appears to be the most adequate solution: Data practices should, one way or another, reflect their true costs, and privacy externalities should therefore be internalised. From the three strategies considered, making individuals accountable for the externality generated by their behaviour — i.e. for their damage of the privacy of others[74] — appears to be the best way to address the phenomenon. While individuals would have a duty not to share third-parties' data without their consent, data controllers would be responsible for not requesting it. These duties of care would then be supplemented by three things:

- a certain degree of technology, to ease the burden (without, however, miraculously making it go away completely),
- a certain regulatory oversight regarding the uses of personal data that should not be allowed in the first place, because necessarily too harmful to others, and

[74] This would be similar to the fact that, e.g. the insurance of reckless drivers does not cover them when it can be shown that their negligence led to the damage.

- a certain adaptation in our practices — those practices which remain too problematic despite the increased accountability and technology, and which are too complicated to regulate and/or too beneficial.

This blend moreover kills two birds with one stone, insofar as it would appropriately address both the issue of privacy externalities which the N&C model alone cannot solve without disproportional side-effects and the issue of burdensomeness for which this model is criticised in the first place. The suggestion I make thus essentially amounts to upholding N&C while complementing it. While I have been critical of the N&C model of privacy-protection, I do not deny the importance of individual control — I only argue that, unaided, it cannot solve the issue of privacy externalities.

Chart 1 below compares graphically the possibilities I have examined to protect privacy adequately: N&C (alone), strategy 1 (tech fixes), 2 (government taking care of privacy) and 3 (duty of care), as well as the blend I am suggesting. The factors are:

- 'how much in control one (really) is of one's privacy, in this model.'
- 'how disproportionate this model — and its side-effects — is if adopted only for the sake of addressing privacy externalities' (i.e. disproportionate compared to the issue of having only imperfect privacy-protection). [75]
- 'how burdensome it is for the individual to achieve full privacy-protection in this model' (i.e. including privacy externalities).

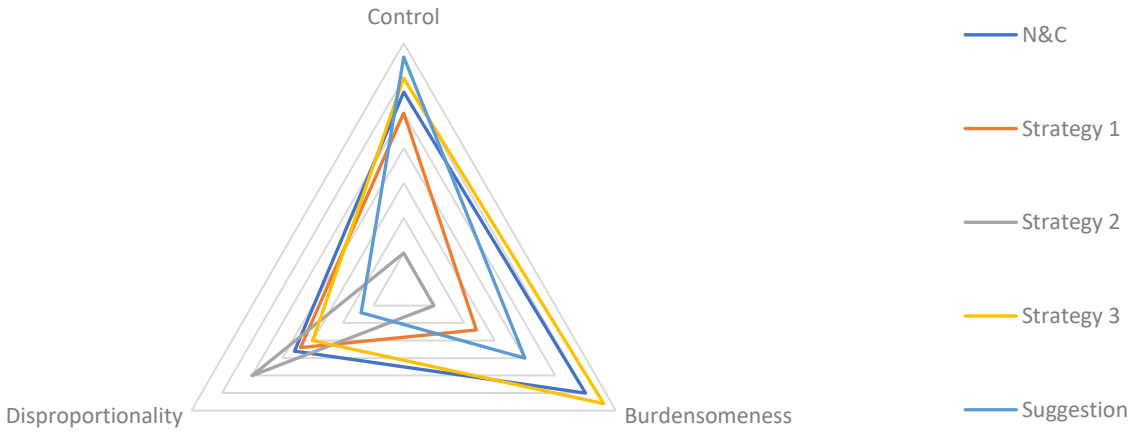

Chart 1: Graphic Comparison of Possible Models of Privacy Protection

[75] Analogously, what could also be taken into account is 'how demanding of a change it would be to switch from N&C to this model.'

While models on this chart are subjectively gauged (since the factors cannot be quantified and measured precisely), this chart is only intended to communicate visually the superiority which I take my suggestion to have, as it appears to secure the value of control while faring better than any of the other models alone.

Further research is required to reinforce the suggestion I make. As a matter of fact, it is not detailed enough yet to be anything more than a promising lead, and I will elaborate below on the points for which research is particularly needed. However, I first want to highlight to anyone willing to carry my suggestion further that questioning the societal, legal and philosophical value of privacy itself was out of the scope of this thesis, and hence that my suggested solution was made under the assumption that privacy matters. This is another reason why it is preferred to, say, a deflationary strategy whereby we simply loosen our concern for privacy. That assumption needs to be accounted for if we raise the issue of privacy externalities to a debate, for the latter may represent less of an issue to those who do not take the value of privacy for granted — especially in the face of the efforts required to consistently protect its.

To those willing to elaborate on my suggestions, further work is needed to determine, among other things, the extent to which one would be responsible for privacy externalities, the kind of legal (or merely social) expectations of due care, the threshold at which one becomes accountable, and the extent to which our convenient and sometimes necessary practices need to be amended. Inspiration could plausibly come from the legal application of the duty of care between private individuals — especially in the context of insurances (e.g. Peck 1960) and of bioethics (Weaver 2016), I believe. It could also come from the field of economics, where appropriate methods have been developed to address externalities,[76] and where game theory is strongly developed.[77] Other promising paths include Helen Nissenbaum's work on privacy as contextual integrity (Nissenbaum 2004),[78] existing work on the impact of 'privacy harms,'[79] as well as the eight criteria examined by the Supreme Court of Tenessee in the due care case of

[76] Cf. Laudon 1996.
[77] Cf. Biczok and Chia 2013 and their Interdependent Privacy Game (regarding Facebook apps), and Humbert *et al.*'s (2015) game-theoretic approach to genetic privacy (in family setting).
[78] Cf. Hull *et al.* (2011) who specifically apply Nissenbaum's work to interdependent cases of privacy. Cf. also Manson and O'Neill's focus on obligations of confidentiality, which "can make clear demands even where there is no explicit professional or contractual relationship" (i.e. even between private individuals), as well as on "systems of accountability" (2008:124-5).
[79] Foreseeably, relevant harm could include monetary or reputational harm, whether definite or probabilistic (i.e. risk), time (lost dealing with the problem), or harm less tangible such as loss of dignity and emotional distress. Importantly, given important losses of privacy be distributed in inconspicuous small doses for multiple individuals, aggregated harm, as well as harm harm to society at large, would also be taken into consideration (Nisenbaum 2010:242-3). See MacCarthy 2011:59-62; Redden and Brand 2017.

1995 (*McCall v. Wilder*, 913 S.W.2d 150, 153).[80] To assess whether due care was taken, the court examined:

- the foreseeability of the harm or injury;
- the possible magnitude of the potential harm or injury;
- the importance or social value of the activity engaged in by the defendant;
- the usefulness of the conduct to the defendant;
- the feasibility of alternative conduct;
- the costs and burdens associated with the alternative conduct;
- the relative usefulness of the alternative conduct;
- the relative safety of the alternative conduct.

These criteria could help determine to what extent one should be held accountable for the damage negligently brought on others' privacy, and the proportionality of the measures that could have prevented it. Moreover, Nissenbaum's norms of "appropriateness" and of "distribution" would help determining in more detail in which contexts the duty of care really matters to the point that negligence should be sanctioned[81] (Nissenbaum 2004:120, 122). However, this should probably not be set in stone, as social expectations and norms of appropriateness evolve — a prime case (though one in which the evolution was forced) being the expectations of privacy in a world of ubiquitous photography, compared to when the technology appeared.

Finally, the solution I propose should plausibly be coupled with MacCarthy's (2011) solution of 'unfairness framework,' which remains today the most comprehensive account of the issue of privacy externalities (though MacCarthy's concept of the latter only concerns category (4) of the four categories of interdependent privacy discussed in chapter 3 (2011:25) — the category on which I focused the least). MacCarthy's unfairness framework reflects strategy 2 which says that data use, rather than just its collection, should be regulated. This framework is intended to guide policy-makers in distinguishing the kind of uses of personal data for which the data subject's consent is relevant, from those uses where consent is not relevant at all (i.e. uses which are so socially harmful that they are impermissible, and uses that are so important to society that one's lack of consent is irrelevant).[82]

[80] MacCarthy (2011) has also established a framework (of "unfairness) which takes into consideration most of the criteria used in the court's case on due care, which he based on the FTC' under-used Unfairness Authority.
[81] For example, a preliminary reading would suggest that one should not be accountable for externalities of networked privacy (where information about the user can be inferred from her network through the principle of homophily), to the contrary of uploading pictures that include recognizable people without their permission, or sharing one's genetic data with commercial parties without consulting one's family beforehand.
[82] A similar framework has been proposed by Cate (2006).

There is thus considerable work left for further research. However, despite the work remaining to be done, if my analysis of the issue at hand is correct, my contribution to the philosophical, societal and legal debate on privacy could prove very valuable. This would be the case insofar as I raised an issue hitherto not noticed in its entirety which, ultimately, threatens the very possibility of privacy self-management, and insofar as the way I analyse and frame this issue clarifies in what sense it is an issue at all, as well as how it should be addressed.

6. Conclusion

The concept of privacy, how it is protected, and the practices that surround it, are recurrently being scrutinised by academics, regulators and the industry, as new practices appear. However, most of the debate has hitherto focused on individuals (what they need to do, what their rights are, the impact that certain practices have on them, etc.) and data controllers (what their duties are, what they actually do behind the scenes, etc.). Privacy protection is often reduced to a relationship between these two actors, and the only other third-party actors discussed — if at all — are commercial (e.g. data brokers) or governmental (e.g. the police or regulators).

In contrast, in this thesis I have tried to show that other actors — private individuals — should also be included in the picture, because appropriate privacy-protection cannot be achieved without taking them into account. Because the impact that these third-party subjects have on one's privacy is currently not addressed, one cannot realistically attain the ideal of privacy self-management — which is already threatened (insofar as it is supposed to be a fundamental right available to all) by the burden of self-management.

Privacy self-management is currently only an ideal, not a reality. Regardless of how someone copes with the burden of privacy self-management (controlling how much one discloses about oneself to others), one has little control over how much others disclose about oneself. This is because in some aspects, privacy is fundamentally an interdependent matter, whereby privacy externalities can be imposed on others by an individual; this is something which is has often been overlooked, and which has therefore been left unregulated. Following the historical development and conceptualisation of information privacy, the most logical way to govern this 'new' aspect of privacy would be to increase and expand measures of N&C. This, however, would greatly add to the aforementioned burden, and would only be possible through the use of algorithms, due to the scale of the phenomenon (i.e. as digital material is uploaded at a rate impossible for humans to manually supervise). Algorithms of this kind (especially those to detect copyright violations) are being criticised for being imperfect and having negative societal effects (mainly on free speech), and could thus have disproportionate side-effects.

I tried to show the magnitude of the phenomenon of interdependent privacy, addressing at length two cases in particular (contact lists and pictures). Together with the other cases raised, these showed how diversely the issue can manifest itself, and the kind of impact it can have. I termed this impact on third-party subjects' privacy an 'externality,' already having a solution in mind to internalise it. To arrive to that solution, I first explored three models of privacy-protection alternative or complementary to N&C. The suggested solution itself took

the best of each world (i.e. strategy), which turned out to have each three a different propensity to address either the 'burdensomeness' or the 'disproportionality' issues.

Therefore, I can now answer my research question, which was *"To what extent can the N&C model of privacy-protection address, unaided, the issue of privacy externalities, and what could (or should) supplement it in this task?"* Although the framework of N&C can to a certain extent be applied to the reality of interdependent privacy, to do so would only add to a burden of privacy management which is already too high, and could moreover have disproportionate consequences. A plausible solution, if we are to address the interdependent aspect of privacy and its invisible costs (externalities), is to complement Notice-and-Consent with a framework of individual accountability, coupled with regulation and technology aimed at reasonably alleviating the burden, while adapting the practices which cannot be solved by those three to the demands of privacy. N&C can remain a major tenet of privacy-protection, but it should not be the only one; while N&C is crucial for the 'collection' aspect of data, accountability should complement it to address the 'use' aspect. Although the solution I suggest is only a rough sketch and a lot is left to further research to carve out the details, this direction is very promising.

The main contribution this thesis makes to the (societal, philosophical and legal) debate about privacy is therefore threefold. Firstly, it raises the issue at hand, i.e. that we cannot meaningfully be said to have privacy self-management if others continuously, inadvertently and in very different ways damage our privacy.[83] Secondly, it accurately frames this issue so that it reflects its complexity, i.e. the generation of externalities due to a failure to perform one's duty of care for others. Thirdly, it suggests and sketches the direction we need to follow to adequately respond to it, i.e. socially and legally expecting from people that they perform this duty of care, easing the burden of privacy management, and amending some of our practices.

This thesis intended to engage in the debate about privacy in a way that encourages reflection about the very foundations of the concept and the way it has hitherto been protected, and it has achieved this goal. The N&C model is deeply limited, and therefore inadequate to ensure real privacy-protection for everyone — not only because it currently overlooks a wide range of stakeholders in the way in which it is applied (a problem of enforcement) but also because, fundamentally, it cannot solve the problem of privacy externalities unaided. The concept of privacy-as-control, the expectation of privacy self-management, and the way these are implemented through a model of privacy-protection that mostly relies on N&C should be

[83] This thesis also incidentally serves as a compilation and generalisation of most of the literature on the topic, which hitherto remained scattered and specific to one application or context.

thoroughly questioned. Under the light brought about by this thesis, the very idea and expectation of self-management become problematic, as they overlook an aspect of the individual's privacy that necessarily depends on others and which allows for the creation of privacy externalities. Privacy self-management as it currently is implemented does not, and cannot alone, provide everyone with meaningful control over their data.

While this conclusion is tailored to the context of privacy protection, a similar narrative could be perhaps be applied to other contexts where the liberal, individualistic framework that is prevalent in the West has imbued the individual with a number of rights — and duties — to self-management. If it is assumed in those cases that individual freedom (conceived as control) alone allows for better protection than other (complementary or alternative) means, that assumption is seriously challenged by the criticism of privacy self-management voiced here.

7. Bibliography

Aldhous, Peter. 'We Tried To Find 10 BuzzFeed Employees Just Like Cops Did For The Golden State Killer'. BuzzFeed, 9 April 2019. https://www.buzzfeednews.com/article/peteraldhous/golden-state-killer-dna-experiment-genetic-genealogy.

Alge, Bradley J., Gary A. Ballinger, Subrahmaniam Tangirala, and James L. Oakley. 'Information Privacy in Organizations: Empowering Creative and Extrarole Performance.' Journal of Applied Psychology 91, no. 1 (2006): 221–32. https://doi.org/10.1037/0021-9010.91.1.221.

Allen, Anita L. Unpopular Privacy: What Must We Hide? Studies in Feminist Philosophy. New York, N.Y: Oxford University Press, 2011.

Amazon. 'Announcing Alexa for Business: Using Amazon Alexa's Voice Enabled Devices for Workplaces'. Amazon Web Services, 30 November 2017. https://aws.amazon.com/blogs/aws/launch-announcing-alexa-for-business-using-amazon-alexas-voice-enabled-devices-for-workplaces/.

Ang, Li, Li Qinghua, and Gao Wei. 'PrivacyCamera: Cooperative Privacy-Aware Photographing with Mobile Phones'. San Diego, CA, USA, 2017. http://ieeexplore.ieee.org/servlet/opac?punumber=7963990.

Anton, Annie, Julia Earp, Qingfeng He, William Stufflebeam, Davide Bolchini, and Carlos Jensen. 'Financial Privacy Policies and the Need for Standardization'. IEEE Security & Privacy 2, no. 36 (2004): 42–44.

Article 29 Data Protection Working Party. 'Opinion 02/2012 on Facial Recognition in Online and Mobile Services'. Brussels, a 2012. https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2012/wp192_en.pdf.

———. 'Opinion 3/2012 on Developments in Biometric Technologies'. Brussels, b 2012. https://ec.europa.eu/justice/article-29/documentation/opinion-recommendation/files/2012/wp193_en.pdf.

Baeles, J. Howard. 'The FTC's Use of Unfairness Authority: Its Rise, Fall, and Resurrection'. Journal of Public Policy & Marketing 192 (2003).

Barth, Susanne, and Menno D.T. de Jong. 'The Privacy Paradox – Investigating Discrepancies between Expressed Privacy Concerns and Actual Online Behavior – A Systematic

Literature Review'. Telematics and Informatics 34, no. 7 (November 2017): 1038–58. https://doi.org/10.1016/j.tele.2017.04.013.

Bernal, Paul. Internet Privacy Rights: Rights to Protect Autonomy. Cambridge: Cambridge University Press, 2014. https://doi.org/10.1017/CBO9781107337428.

Besmer, Andrew, and Heather Richter Lipford. 'Users' (Mis)Conceptions of Social Applications'. Proceedings of Graphics Interface 2010, 2010, 63–70.

Biczók, Gergely, and Pern Hui Chia. 'Interdependent Privacy: Let Me Share Your Data'. In Financial Cryptography and Data Security, edited by Ahmad-Reza Sadeghi, 7859:338–53. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013. https://doi.org/10.1007/978-3-642-39884-1_29.

Big Brother Watch. 'Stop the Met Police Using Authoritarian Facial Recognition Cameras'. CrowdJustice, April 2019. https://www.crowdjustice.com/case/face-off/.

Bloustein, E. J. Individual and Group Privacy. Transaction Publishers. New Brunswick, 1978.

Bowie, Norman E., and Karim Jamal. 'Privacy Rights on the Internet: Self-Regulation or Government Regulation?' Business Ethics Quarterly 16, no. 3 (July 2006): 323–42. https://doi.org/10.5840/beq200616340.

Brignull, Harry. 'Dark Patterns'. Dark Patterns, 2019. https://www.darkpatterns.org/.

Brown, Kristen. 'A Researcher Needed Three Hours to Identify Me From My DNA'. Bloomberg, 12 April 2019. https://www.bloomberg.com/news/articles/2019-04-12/a-researcher-needed-three-hours-to-identify-me-from-my-dna.

Buolamwini, Joy, and Timnit Gebru. 'Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification'. Proceedings of Machine Learning Research, Conference on Fairness, Accountability, and Transparency, 81, no. 1 (2018): 1–15.

Burt, Andrew, and Dan Geer. 'The End of Privacy'. The New York Times, 5 October 2017, sec. Opinion. https://www.nytimes.com/2017/10/05/opinion/privacy-rights-security-breaches.html.

Burt, Chris. 'Driver's License Photos Increasingly Used by U.S. Police for Identification with Facial Recognition'. Biometric Update, 18 June 2018. https://www.biometricupdate.com/201806/drivers-license-photos-increasingly-used-by-u-s-police-for-identification-with-facial-recognition.

———. ‘IBM Launches Public Data Set to Further Research into Diversity and Facial Biometrics’. Biometric Update, 2019. https://www.biometricupdate.com/201901/ibm-launches-public-data-set-to-further-research-into-diversity-and-facial-biometrics.

Buttarelli, Giovanni, and Data Protection Officers of European Community institutions and bodies. ‘The EDPS Video-Surveillance Guidelines’. Bruxelles: European Data Protection Supervisor (EDPS), 7 July 2009. https://edps.europa.eu/sites/edp/files/publication/09-07-07_consultation_videosurveillance_guidelines_en.pdf.

Calo, Ryan. ‘M. Ryan Calo, Against Notice Skepticism in Privacy (and Elsewhere)’. Notre Dame Law Review 1027, no. 1033 (2012).

Cate, Fred H. ‘The Failure of Fair Information Practice Principles’. In Consumer Protection in the Age of the Information Economy, edited by Jane K. Winn. Markets and the Law. Aldershot, Hants, England ; Burlington, VT, 2006.

Caughlin, T. Trevor, Nick Ruktanonchai, Miguel A. Acevedo, Kenneth K. Lopiano, Olivia Prosper, Nathan Eagle, and Andrew J. Tatem. ‘Place-Based Attributes Predict Community Membership in a Mobile Phone Communication Network’. Edited by Angel Sánchez. PLoS ONE 8, no. 2 (22 February 2013): e56057. https://doi.org/10.1371/journal.pone.0056057.

Center for Democracy and Technology. ‘CDT Deferal Baseline Privacy Legislation Discussion Draft’. Center for Democracy and Technology (CDT), 2018. https://cdt.org/campaign/federal-privacy-legislation/.

Chadwick, Ruth F., Mairi Levitt, and Darren Shickle, eds. The Right to Know and the Right Not to Know: Genetic Privacy and Responsibility. Second edition. Cambridge Bioethics and Law. Cambridge, United Kingdom: Cambridge University Press, 2014.

Chaudhry, Amir, Jon Crowcroft, Heidi Howard, Anil Madhavapeddy, Richard Mortier, Hamed Haddadi, and Derek McAuley. ‘Personal Data: Thinking Inside the Box’. Aarhus Series on Human Centered Computing 1, no. 1 (5 October 2015): 4. https://doi.org/10.7146/aahcc.v1i1.21312.

Chen, Angela. ‘The Guy Who Made a Tool to Track Women in Porn Videos Is Sorry’. MIT Technology Review, 31 May 2019. https://www.technologyreview.com/s/613607/facial-recognition-porn-database-privacy-gdpr-data-collection-policy/.

Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim. ‘Privacy and Personal Data Collection with Information Externalities’. Journal of Public Economics 173 (May 2019): 113–24. https://doi.org/10.1016/j.jpubeco.2019.02.001.

Cohen, Julie. ‘Configuring the Networked Citizen’. In Imagining New Legalities: Privacy and Its Possibilities in the 21st Century, edited by Austin Sarat et al., 2012.

Conger, Kate, Richard Fausset, and Serge F. Kovaleski. ‘San Francisco Bans Facial Recognition Technology’. The New York Times, 16 May 2019, sec. U.S. https://www.nytimes.com/2019/05/14/us/facial-recognition-ban-san-francisco.html.

Conly, Sarah. ‘Against Autonomy: Justifying Coercive Paternalism’. Journal of Medical Ethics 40, no. 5 (2013): 349–349. https://doi.org/10.1136/medethics-2013-101444.

Constine, Josh. ‘Facebook Will Reveal Who Uploaded Your Contact Info for Ad Targeting’. TechCrunch (blog), February 2019. http://social.techcrunch.com/2019/02/06/why-am-i-seeing-this-ad/.

‘Cozy Cloud’. Cozy Cloud Website, 2017. https:// cozy.io/.

Crandall, D. J., L. Backstrom, D. Cosley, S. Suri, D. Huttenlocher, and J. Kleinberg. ‘Inferring Social Ties from Geographic Coincidences’. Proceedings of the National Academy of Sciences 107, no. 52 (28 December 2010): 22436–41. https://doi.org/10.1073/pnas.1006155107.

Cranor, Lorrie Faith. Web Privacy with P3P: The Platform for Privacy Preferences. 1. ed. Beijing: O’Reilly, 2002.

Culnan, Mary J. ‘Protecting Privacy Online: Is Self-Regulation Working?’ Journal of Public Policy & Marketing 19, no. 1 (April 2000): 20–26. https://doi.org/10.1509/jppm.19.1.20.16944.

Culnan, Mary J., and Pamela K. Armstrong. ‘Information Privacy Concerns, Procedural Fairness, and Impersonal Trust: An Empirical Investigation’. Organization Science 10, no. 1 (February 1999): 104–15. https://doi.org/10.1287/orsc.10.1.104.

Cyphers, Benett, and Jason Kelley. ‘Facebook Got Caught Phishing For Friends’. Eletronic Frontier Foundation, 4 April 2019. https://www.eff.org/deeplinks/2019/04/facebook-got-caught-phishing-friends.

danah, boyd. ‘Networked Privacy’. Personal Democracy Forum, New York, 6 June 2011. https://www.danah.org/papers/talks/2011/PDF2011.html.

De Hert, Paul. ‘Identity Management of E-ID, Privacy and Security in Europe. A Human Rights View’. Information Security Technical Report 13, no. 2 (May 2008): 71–75. https://doi.org/10.1016/j.istr.2008.07.001.

Dent, Steve. ‘Amazon Shareholders Will Vote to Ban Facial Recognition Tech’. Engadget, 15 April 2019. https://www.engadget.com/2019/04/15/amazon-shareholder-vote-facial-recognition/.

‘Digi.Me’. Digi.me Website, 2017. https://digi.me.

Doctorow, Cory. ‘The European Copyright Directive: What Is It, and Why Has It Drawn More Controversy Than Any Other Directive In EU History?’ Electronic Frontier Foundation, 19 March 2019. https://www.eff.org/deeplinks/2019/03/european-copyright-directive-what-it-and-why-has-it-drawn-more-controversy-any.

Duch-Brown, Nestor, Bertin Martens, and Frank Mueller-Langer. ‘The Economics of Ownership, Access and Trade in Digital Data’. JRC Digital Economy Working Paper, 2017. https://doi.org/10.2139/ssrn.2914144.

Duhigg, Charles. ‘What Does Your Credit-Card Company Know About You?’ The New York Times, 12 May 2009, sec. Magazine. https://www.nytimes.com/2009/05/17/magazine/17credit-t.html.

EDPS Ethics Advisory Group, J. Peter Burgess, Luciano Floridi, Aurélie Pols, and Jeroen van den Hoven. ‘Towards a Digital Ethics’. European Data Protection Supervisor (EDPS), 2018. https://edps.europa.eu/sites/edp/files/publication/18-01-25_eag_report_en.pdf.

Electronic Frontier Foundation. ‘Fix It Already’. Electronic Frontier Foundation, March 2019. https://fixitalready.eff.org/facebook.

———. ‘Success Story: Dismantling UK’s Biometric ID Database’. Electronic Frontier Foundation, 6 December 2012. https://www.eff.org/pages/success-story-dismantling-uk%E2%80%99s-biometric-id-database.

Endedijk, Bram, and Eric van den Berg. ‘Nederlandse politie wil gebruikmaken van particuliere DNA-databank in VS’. NRC, 5 February 2019. https://www.nrc.nl/nieuws/2019/02/05/politie-kan-bijna-iedereen-vinden-met-particuliere-dna-databank-a3653045.

Enserink, M., and G. Chin. ‘The End of Privacy’. Science 347, no. 6221 (30 January 2015): 490–91. https://doi.org/10.1126/science.347.6221.490.

Erlich, Yaniv, Tal Shor, Itsik Pe'er, and Shai Carmi. 'Identity Inference of Genomic Data Using Long-Range Familial Searches'. Science 362, no. 6415 (9 November 2018): 690–94. https://doi.org/10.1126/science.aau4832.

European Commission. 'An Emerging Offer of "Personal Information Management Services" – Current State of Service Offers and Challenges'. European Commission Report, 2016.

European Data Protection Supervisor. 'EDPS Opinion on Personal Information Management Systems. Towards More User Empowerment in Managing and Processing Personal Data', 2016.

European Digital Rights (EDRi). 'Censorship Machine: Busting the Myths', 13 December 2017. https://edri.org/censorship-machine-busting-myths/.

———. 'Copyright Negotiations Begin to Derail'. EDRi (blog), 21 January 2019. https://edri.org/copyright-negotiations-begin-derail/.

———. 'RE-Deconstructing the Article 13 of the Copyright Proposal as Amended by JURI Committee (Revision 3)', 2018. https://edri.org/files/copyright/20180626-ReDeconstructing_Article13.pdf.

European Parliament, and Council of the European Union. 'Directive 95/46/EC of the European Parliament and of the Council on the Protection of Individuals with Regard to the Processing of Personal Data and on the Free Movement of Such Data'. L281, 24 October 1995.

———. 'Regulation on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data (GDPR), and Repealing Directive 95/46/EC (Data Protection Directive)'. L119, 14 April 2016.

Facebook Inc. 'Facebook Post-Hearing Responses to Commerce Committee: "Facebook, Social Media Privacy, and the Use and Abuse of Data"', June 2018. https://www.judiciary.senate.gov/imo/media/doc/Zuckerberg%20Responses%20to%20Commerce%20Committee%20QFRs1.pdf.

Federal Trade Commission (FTC). 'Facing Facts: Best Practices for Common Uses of Facial Recognition Technologies'. FTC, October 2012.

———. 'Privacy Online: A Report to Congress'. Washington DC: Federal Trade Commission, 1998.

Ferrer, Vincent. ‘Right This Way: A Potential Artificial Intelligence-Based Solution for Complying with Article 13 of the EU’s 2018 Copyright Directive’. Law School Student Scholarship, no. 948 (2019). https://scholarship.shu.edu/student_scholarship/948.

Floridi, Luciano. ‘Soft Ethics, the Governance of the Digital and the General Data Protection Regulation’. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences 376, no. 2133 (28 November 2018): 20180081. https://doi.org/10.1098/rsta.2018.0081.

Fox, Jenna Tesse. ‘With Alexa for Hospitality, Marriott Adds Amazon to the Guest Experience’. Hotel Management, 19 June 2018. https://www.hotelmanagement.net/tech/alexa-for-hospitality-marriott-adds-amazon-to-guest-experience-0.

Gallagher, Andrew C., and Chen Tsuhan. ‘Clothing Cosegmentation for Recognizing People’. In 2008 IEEE Conference on Computer Vision and Pattern Recognition, 1–8. Anchorage, AK, USA: IEEE, 2008. https://doi.org/10.1109/CVPR.2008.4587481.

Garcia, David. ‘Leaking Privacy and Shadow Profiles in Online Social Networks’. Science Advances 3, no. 8 (2017): e1701172. https://doi.org/10.1126/sciadv.1701172.

Garcia-Murillo, Martha, and Ian MacInnes. ‘Così Fan Tutte : A Better Approach than the Right to Be Forgotten’. Telecommunications Policy 42, no. 3 (April 2018): 227–40. https://doi.org/10.1016/j.telpol.2017.12.003.

Gebhart, Gennie. ‘You Gave Facebook Your Number For Security. They Used It For Ads.’ Electronic Frontier Foundation, 27 September 2018. https://www.eff.org/deeplinks/2018/09/you-gave-facebook-your-number-security-they-used-it-ads.

Gellman, Robert. ‘Fair Information Practices: A Basic History’. SSRN Electronic Journal, 2014. https://doi.org/10.2139/ssrn.2415020.

Gras, Marianne L. ‘The Legal Regulation of CCTV in Europe’ CCTV Special (eds. Norris, McCahill and Wood), no. 2 (2004): 216–29. http://www.surveillance-and-society.org/cctv.htm.

Guarino, Ben. ‘Russia’s New FindFace App Identifies Strangers in a Crowd with 70 Percent Accuracy - The Washington Post’. The Washington Post, 18 May 2016. https://www.washingtonpost.com/news/morning-mix/wp/2016/05/18/russias-new-

findface-app-identifies-strangers-in-a-crowd-with-70-percent-accuracy/?utm_term=.0562fd0c85cc.

Hallinan, Dara, and Paul de Hert. ‘Genetic Classes and Genetic Categories: Protecting Genetic Groups Through Data Protection Law’. In Group Privacy, edited by Linnet Taylor, Luciano Floridi, and Bart van der Sloot, 175–96. Cham: Springer International Publishing, 2017. https://doi.org/10.1007/978-3-319-46608-8_10.

Hann, I-H, K-L Hui, TS Lee, and IPL Png. ‘Online Informationprivacy: Measuring the Cost-Benefit Trade-Off’. Twenty-Third International Conference on Information Systems, 2002.

Henn, Steve. ‘Clever Hacks Give Google Glass Many Unintended Powers’. NPR.org, 17 July 2013. https://www.npr.org/sections/alltechconsidered/2013/07/17/202725167/clever-hacks-give-google-glass-many-unintended-powers.

Hill, Kashmir. ‘Facebook Is Giving Advertisers Access to Your Shadow Contact Information’. Gizmodo, 26 September 2018. https://gizmodo.com/facebook-is-giving-advertisers-access-to-your-shadow-co-1828476051.

———. ‘How Facebook Figures Out Everyone You’ve Ever Met’. Gizmodo, 7 November 2017. https://gizmodo.com/how-facebook-figures-out-everyone-youve-ever-met-1819822691.

Hiller, Janine S., and France Belanger. ‘Privacy Strategies for Electronic Government’. In E-Government 2001, edited by Mark Abramson and Grady Means. IBM Endowment for The Business of Government. New York: Rowman and Littlefields Publishers, 2001.

Hooker, Brad. ‘The Demandingness Objection’. In The Problem of Moral Demandingness, edited by T. Chapepell, 148–62. London: Palgrave Macmillan, 2009.

‘Hub of All Things’. Hub of All Things GitHub page, 2017. https://github.com/Hub-of-all-Things.

Hull, Gordon. ‘Successful Failure: What Foucault Can Teach Us about Privacy Self-Management in a World of Facebook and Big Data’. Ethics and Information Technology 17, no. 2 (June 2015): 89–101. https://doi.org/10.1007/s10676-015-9363-z.

Hull, Gordon, Heather Richter Lipford, and Celine Latulipe. ‘Contextual Gaps: Privacy Issues on Facebook’. Ethics and Information Technology 13, no. 4 (December 2011): 289–302. https://doi.org/10.1007/s10676-010-9224-8.

Humbert, Mathias, Erman Ayday, Jean-Pierre Hubaux, and Amalio Telenti. ‘On Non-Cooperative Genomic Privacy’. In Financial Cryptography and Data Security, edited by Rainer Böhme and Tatsuaki Okamoto, 8975:407–26. Berlin, Heidelberg: Springer Berlin Heidelberg, 2015. https://doi.org/10.1007/978-3-662-47854-7_24.

Information Commissioner’s Office (ICO). ‘In the Picture: A Data Protection Code of Practice for Surveillance Cameras and Personal Information’, 2017. https://ico.org.uk/media/for-organisations/documents/1542/cctv-code-of-practice.pdf.

Innes, M. ‘“Control Creep”’. Sociological Research Online 6, no. 3 (2001). http://www.socresonline.org.uk/6/3/innes.html.

Introna, Lucas D. ‘Privacy and the Computer: Why We Need Privacy in the Information Society’. Metaphilosophy 28, no. 3 (July 1997): 259–75. https://doi.org/10.1111/1467-9973.00055.

Jernigan, Carter, and Behram F.T. Mistree. ‘Gaydar: Facebook Friendships Expose Sexual Orientation’. First Monday 14, no. 10 (25 September 2009). https://doi.org/10.5210/fm.v14i10.2611.

Joseph, George. ‘Inside the Surveillance Program IBM Built for Rodrigo Duterte - Type Investigations’. Type Investigations, 20 March 2019. https://www.typeinvestigations.org/investigation/2019/03/20/inside-the-video-surveillance-program-ibm-built-for-philippine-strongman-rodrigo-duterte/.

Kelley, Jason. ‘Shareholders Demand To Know How Northrop Grumman Will Protect Human Rights While Building Massive DHS Database’. Electronic Frontier Foundation, 8 May 2019. https://www.eff.org/deeplinks/2019/05/shareholders-northrop-grumman-demand-know-how-company-will-protect-human-rights.

Kitchin, Rob. The Data Revolution: Big Data, Open Data, Data Infrastructures & Their Consequences. Los Angeles, California: SAGE Publications, 2014a.

———. ‘The Real-Time City? Big Data and Smart Urbanism’. GeoJournal 79, no. 1 (2014b): 1–14. https://doi.org/10.1007/s10708-013-9516-8.

Kokolakis, Spyros. ‘Privacy Attitudes and Privacy Behaviour: A Review of Current Research on the Privacy Paradox Phenomenon’. Computers & Security 64 (January 2017): 122–34. https://doi.org/10.1016/j.cose.2015.07.002.

Koops, B.-J. ‘The Trouble with European Data Protection Law’. International Data Privacy Law 4, no. 4 (1 November 2014): 250–61. https://doi.org/10.1093/idpl/ipu023.

Kupfer, Joseph. ‘Privacy, Autonomy, and Self-Concept’. American Philosophical Quarterly 24, no. 1 (1987): 81–89. https://www.jstor.org/stable/20014176.

Lampinen, Airi. ‘Networked Privacy Beyond the Individual: Four Perspectives to “Sharing”’. Aarhus Series on Human Centered Computing 1, no. 1 (5 October 2015): 4. https://doi.org/10.7146/aahcc.v1i1.21300.

———. ‘Networked Privacy Beyond the Individual: Four Perspectives to “Sharing”’. Aarhus Series on Human Centered Computing 1, no. 1 (5 October 2015): 4. https://doi.org/10.7146/aahcc.v1i1.21300.

Lampinen, Airi, Vilma Lehtinen, Asko Lehmuskallio, and Sakari Tamminen. ‘We’re in It Together: Interpersonal Management of Disclosure in Social Network Services’. edited by Association for Computing Machinery, 2011.

Landreau, Isabelle, Gerard Peliks, Nicolas Binctin, Virginie Pez-Perard, and Leger Lucas. ‘My Data Are Mine: Why We Should Have Ownership Rights on Our Personal Data’. GenerationLibre Think Tank, April 2018.

Laudon, Kenneth C. ‘Markets and Privacy’. Communications of the ACM 39, no. 9 (1996): 92–104. https://doi.org/10.1145/234215.234476.

Legifrance. Cour d’appel de Versailles, 9 février 2017, 15/08667, No. 15/08667 (Cour d’appel de Versailles février 2017).

———. Cour de cassation, civile, Chambre civile 1, 6 octobre 2011, 10-23.606, Inédit, Inédit (Cour de cassation 2011).

Lehtiniemi, Tuukka, and Yki Kortesniemi. ‘Can the Obstacles to Privacy Self-Management Be Overcome? Exploring the Consent Intermediary Approach’. Big Data & Society 4, no. 2 (December 2017): 205395171772193. https://doi.org/10.1177/2053951717721935.

Liao, Shannon. ‘IBM Didn’t Inform People When It Used Their Flickr Photos for Facial Recognition Training’. The Verge, 12 March 2019. https://www.theverge.com/2019/3/12/18262646/ibm-didnt-inform-people-when-it-used-their-flickr-photos-for-facial-recognition-training.

Litman-Navarro, Kevin. ‘We Read 150 Privacy Policies. They Were an Incomprehensible Disaster.’ The New York Times, 12 June 2019, sec. Opinion. https://www.nytimes.com/interactive/2019/06/12/opinion/facebook-google-privacy-policies.html, https://www.nytimes.com/interactive/2019/06/12/opinion/facebook-google-privacy-policies.html.

Livingstone, Rob. ‘Smile! Face Recognition for Google Glass Is Here, Thanks to Hackers’. The Conversation, 28 July 2013. http://theconversation.com/smile-face-recognition-for-google-glass-is-here-thanks-to-hackers-16262.

Loebel, Jens-Martin. ‘Is Privacy Dead? – An Inquiry into GPS-Based Geolocation and Facial Recognition Systems’. In ICT Critical Infrastructures and Society, edited by Magda David Hercheui, Diane Whitehouse, William McIver, and Jackie Phahlamohlaka, 386:338–48. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012. https://doi.org/10.1007/978-3-642-33332-3_31.

Lohr, Steve. ‘Facial Recognition Is Accurate, If You’re a White Guy’. The New York Times, 8 February 2018, sec. Technology. https://www.nytimes.com/2018/02/09/technology/facial-recognition-race-artificial-intelligence.html.

Lynch, Jennifer. Hearing on Law Enforcement’s Use of Facial Recognition Technology, § United States House Committee on Oversight and Government Reform (2017).

MacCarthy, Mark. ‘New Directions in Privacy: Disclosure, Unfairness and Externalies’. I/S: A Journal of Law and Policy for the Information Society 6 (2011): 425.

Madden, Mary. ‘The Devastating Consequences of Being Poor in the Digital Age’. The New York Times, 25 April 2019.

Madden, Mary, Michelle E. Gilman, Karen Levy, and Alice E Marwick. ‘Privacy, Poverty and Big Data: A Matrix of Vulnerabilities for Poor Americans (March 9, 2017)’. Washington University Law Review 53, no. 95 (2017). https://ssrn.com/abstract=2930247.

Manson, Neil C., and Onora O’Neill. Rethinking Informed Consent in Bioethics. Reprint. Cambridge: Cambridge Univ. Press, 2008.

Marwick, Alice E, and danah boyd. ‘Networked Privacy: How Teenagers Negotiate Context in Social Media’. New Media & Society 16, no. 7 (2014): 1051–67. https://doi.org/10.1177/1461444814543995.

McCALL v. WILDER | 913 S.W.2d 150 (1995) | w2d15011063 (Supreme Court of Tenessee 11 December 1995).

McDaniel, Patrick, and Stephen McLaughlin. ‘Security and Privacy Challenges in the Smart Grid’. IEEE Security & Privacy Magazine 7, no. 3 (May 2009): 75–77. https://doi.org/10.1109/MSP.2009.76.

McDonald, Aleecia M., and Lorrie Faith Cranor. ‘The Cost of Reading Privacy Policies’. I/S: A Journal of Law and Policy for the Information Society 4, no. 3 (2008).

McRobb, Steve, and Simon Rogerson. ‘Are They Really Listening?: An Investigation into Published Online Privacy Policies at the Beginning of the Third Millennium’. Information Technology & People 17, no. 4 (December 2004): 442–61. https://doi.org/10.1108/09593840410570285.

‘Meeco’. Meeco Website, 2017. https://meeco.me.

Milne, George R., and Mary J. Culnan. ‘Strategies for Reducing Online Privacy Risks: Why Consumers Read (or Don’t Read) Online Privacy Notices’. Journal of Interactive Marketing 18, no. 3 (January 2004): 15–29. https://doi.org/10.1002/dir.20009.

Mims, Christopher. ‘Amazon’s Plan to Move In to Your Next Apartment Before You Do’. Wall Street Journal, 1 June 2019, sec. Tech. https://www.wsj.com/articles/amazons-plan-to-move-in-to-your-next-apartment-before-you-do-11559361605.

Mindle, Grant B. ‘Liberalism, Privacy, and Autonomy’. The Journal of Politics 51, no. 3 (August 1989): 575–98. https://doi.org/10.2307/2131496.

Molteni, Megan. ‘The US Urgently Needs New Genetic Privacy Laws’. Wired, 1 May 2019. https://www.wired.com/story/the-us-urgently-needs-new-genetic-privacy-laws/.

Montjoye, Yves-Alexandre de, Erez Shmueli, Samuel S. Wang, and Alex Sandy Pentland. ‘OpenPDS: Protecting the Privacy of Metadata through SafeAnswers’. Edited by Tobias Preis. PLoS ONE 9, no. 7 (9 July 2014): e98790. https://doi.org/10.1371/journal.pone.0098790.

Moore, Adam D. ‘Toward Informational Privacy Rights’. San Diego Law Review 44, no. 809 (2007). https://heinonline.org/HOL/LandingPage?handle=hein.journals/sanlr44&div=43&id=&page=&t=1557070527.

Morozov, E. (2013) To Save Everything, Click Here: Technology, Solutionism, and the Urge to Fix Problems That Don’t Exist. Allen Lane, New York., n.d.

Morozov, Evgeny. The Net Delusion: The Dark Side of Internet Freedom. 1st ed. New York: Public Affairs, 2011.

Newman, Nathan. ‘This Time Is Different: How Big Data Has Left the Middle Class Behind’. Huffington Post, 24 March 2015. https://www.huffpost.com/entry/this-time-is-different-ho_b_6931044.

Nissenbaum, Helen F. ‘A Contextual Approach to Privacy Online’. Daedalus 140, no. 4 (2011): 32–48.

———. ‘Privacy as Contextual Integrity’. Washington Law Review 79, no. 1 (2004): 119–57.

———. Privacy in COntext: Technology, Policy and the Integrity of Social Life. Stanford University Press, 2010.

Obar, Jonathan A. ‘Big Data and The Phantom Public: Walter Lippmann and the Fallacy of Data Privacy Self-Management’. Big Data & Society 2, no. 2 (27 December 2015): 205395171560887. https://doi.org/10.1177/2053951715608876.

Obar, Jonathan A., and Anne Oeldorf-Hirsch. ‘The Biggest Lie on the Internet: Ignoring the Privacy Policies and Terms of Service Policies of Social Networking Services’. Information, Communication & Society, 3 July 2018, 1–20. https://doi.org/10.1080/1369118X.2018.1486870.

Office of the Privacy Commissioner of Canada (OPC). ‘2016–2017 Annual Report to Parliament on the Personal Information Protections and Electronic Documents Act’, 2017.

Organisation for Economic Co-operation and Development (OECD). ‘OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data’, 1980.

Parket bij de Hoge Raad. ‘Arrest Hoge Raad: ECLI:NL:HR:2010:BK6331. Zaaknummer 08/04524 B’, 23 March 2010. http://deeplink.rechtspraak.nl/uitspraak?id=ECLI:NL:PHR:2010:BK6331.

Pascalev, Mario. ‘Privacy Exchanges: Restoring Consent in Privacy Self-Management’. Ethics and Information Technology 19, no. 1 (March 2017): 39–48. https://doi.org/10.1007/s10676-016-9410-4.

Payne, Chris. ‘Everything You Need to Know About CCTV and the GDPR’. Infinigate, 8 January 2018. https://blog.infinigate.co.uk/everything-you-need-to-know-about-cctv-gdpr.

Peck, Cornelius J. ‘Comparative Negligence and Automobile Liability Insurance’. Michigan Law Review 58, no. 5 (March 1960): 689. https://doi.org/10.2307/1285823.

Piattelli--Palmarini, M. Inevitable Illusions: How Mistakes of Reason Rule Our Minds. New York: John Wiley & Sons, 1994.

Piattelli-Palmarini, Massimo. Inevitable Illusions: How Mistakes of Reason Rule Our Minds. New York: Wiley, 1994.

Poikola, A., K. Kuikkaniemi, and H. Honko. 'MyData – A Nordic Model for Human-Centered Personal Data Management and Processing'. Helsinki: Finnish Ministry of Transport and Communications, 2015.

Privacy International. 'Betrayed by an App She Had Never Heard of" - How TrueCaller Is Endangering Journalists'. Privacy International, 2019b. http://www.privacyinternational.org/case-study/2997/betrayed-app-she-had-never-heard-how-truecaller-endangering-journalists.

———. 'Biometrics Know No Borders, They Must Be Subject To Extreme Vetting'. Privacy International, 25 October 2017. http://privacyinternational.org/blog/653/biometrics-know-no-borders-they-must-be-subject-extreme-vetting.

———. 'Most Cookie Banners Are Annoying and Deceptive. This Is Not Consent.' Privacy International, 2019a. http://privacyinternational.org/explainer/2975/most-cookie-banners-are-annoying-and-deceptive-not-consent.

Quain, John. 'How 3,000 Streetlights Turned San Diego into America's Smartest City'. Digital Trends, 25 July 2018. https://www.digitaltrends.com/cool-tech/how-3000-streetlights-turned-san-diego-into-americas-smartest-city/.

Rainie, Lee. 'How Americans Feel about Social Media and Privacy'. Pew Research Center, 27 March 2018. https://www.pewresearch.org/fact-tank/2018/03/27/americans-complicated-feelings-about-social-media-in-an-era-of-privacy-concerns/.

Recuerda, Miguel A. 'Risk and Reason in the European Union Law'. European Food and Feed Law Review 5 (2006).

Redden, Joanna, and Jessica Brand. 'Data Harm Record – Data Justice Lab'. Data Justice Lab, 9 December 2017. https://datajusticelab.org/data-harm-record/.

Reidenberg, Joel, Cameron Russel, Alexander Callen, Sophia Qasir, and Thomas Norton. 'Privacy Harms and the Effectiveness of the Notice and Choice Framework'. A Journal of Law and Policy for the Information Society 11, no. 2 (2014).

Roessler, Beate, and Dorota Mokrosinska. ‘Privacy and Social Interaction’. Philosophy & Social Criticism 39, no. 8 (October 2013): 771–91. https://doi.org/10.1177/0191453713494968.

Sarigol, Emre, David Garcia, and Frank Schweitzer. ‘Online Privacy as a Collective Phenomenon’. In Proceedings of the Second Edition of the ACM Conference on Online Social Networks - COSN ’14, 95–106. Dublin, Ireland: ACM Press, 2014. https://doi.org/10.1145/2660460.2660470.

Schneiders, Bruce. ‘A Taxonomy of Social Networking Data’. IEEE Security & Privacy, 2010.

Schrems, Max. ‘Complaint against Facebook Ireland Ltd. – 02 “Shadow Profiles”’, 18 August 2011. http://www.europe-v-facebook.org/Compalint_02_Shadow_Profiles.pdf.

Smit, Barry. ‘Leuk, een DNA-test. Maar ken je de privacyrisico’s?’ Bits of Freedom, 23 December 2018. https://www.bitsoffreedom.nl/2018/12/23/leuk-een-dna-test-maar-ken-je-de-risicos/.

Solove, Daniel J. ‘Conceptualizing Privacy’. California Law Review 90 (2002).

———. ‘Privacy Self-Management and the Consent Dilemma’. Harvard Law Review 126 (2013).

———. The Future of Reputation: Gossip, Rumor, and Privacy on the Internet. , 2007. Yale University Press, 2007. www.jstor.org/stable/j.ctt1npqjw.

———. Understanding Privacy. Cambridge, Mass: Harvard University Press, 2008.

Solove, Daniel J., Marc Rotenberg, and Paul M. Schwartz. Privacy, Information, and Technology. New York: Aspen Publishers, 2006.

Sonderholm, Jorn. ‘World Poverty, Positive Duties, and the Overdemandingness Objection’. Politics, Philosophy & Economics 12, no. 3 (August 2013): 308–27. https://doi.org/10.1177/1470594X12447779.

Song, Yang, and Thomas Leung. ‘Context-Aided Human Recognition – Clustering’. In Computer Vision – ECCV 2006, edited by Aleš Leonardis, Horst Bischof, and Axel Pinz, 3953:382–95. Berlin, Heidelberg: Springer Berlin Heidelberg, 2006. https://doi.org/10.1007/11744078_30.

Strahilevitz, Lior Jacob. ‘Toward a Positive Theory of Privacy Law’. Harvard Law Review., no. 126 (2010).

Stufflebeam, William H., Annie I. Anton, Qingfeng He, and Neha Jain. ‘Specifying Privacy Policies with P3P and EPAL: Lessons Learned’. In Proceedings of the 2004 ACM Workshop on Privacy in the Electronic Society - WPES ’04, 35. Washington DC, USA: ACM Press, 2004. https://doi.org/10.1145/1029179.1029190.

Symeonidis, Iraklis, Fatemeh Shirazi, Gergely Biczók, Cristina Pérez-Solà, and Bart Preneel. ‘Collateral Damage of Facebook Apps: Friends, Providers, and Privacy Interdependence’. In ICT Systems Security and Privacy Protection, edited by Jaap-Henk Hoepman and Stefan Katzenbeisser, 471:194–208. Cham: Springer International Publishing, 2016. https://doi.org/10.1007/978-3-319-33630-5_14.

Tait, Amelia. ‘Why Does Facebook Recommend Friends I’ve Never Even Met?’ Wired UK, 29 May 2019. https://www.wired.co.uk/article/facebook-people-you-may-know-friend-suggestions.

Tang, Zhulei, Yu (Jeffrey) hu, and Michael D. smith. ‘Gaining Trust Through Online Privacy Protection: Self-Regulation, Mandatory Standards, or Caveat Emptor’. Journal of Management Information Systems 24, no. 4 (April 2008): 153–73. https://doi.org/10.2753/MIS0742-1222240406.

Tavani, Herman T. ‘Genomic Research and Data-Mining Technology: Implications for Personal Privacy and Informed Consent’. Ethics and Information Technology 6, no. 1 (March 2004): 15–28. https://doi.org/10.1023/B:ETIN.0000036156.77169.31.

Taylor, Linnet, Luciano Floridi, and Bart van der Sloot, eds. Group Privacy: New Challenges of Data Technologies. Cham: Springer International Publishing, 2017. https://doi.org/10.1007/978-3-319-46608-8.

———. ‘Introduction: A New Perspective on Privacy’. In Group Privacy, edited by Linnet Taylor, Luciano Floridi, and Bart van der Sloot, 1–12. Cham: Springer International Publishing, 2017. https://doi.org/10.1007/978-3-319-46608-8_1.

Taylor, Simon. ‘Data: The New Currency?’ European Voice, 2017. http://www.politico.eu/wp-content/uploads/2014/10/Data-report.pdf

Trappel, Josef, ed. Digital Media Inequalities: Policies against Divides, Distrust and Discrimination. Nordicom, 2019. https://www.nordicom.gu.se/en/publikationer/digital-media-inequalities.

Vallet, Felicien. ‘Les droits de la voix (1/2) : Quelle écoute pour nos systèmes ?’ Laboratoire d’Innovation Numerique de la CNIL (LINC), 13 May 2019. https://linc.cnil.fr/fr/les-droits-de-la-voix-12-quelle-ecoute-pour-nos-systemes.

Vallet, Felicien, and Jean-François Bonastre. ‘Jean-François Bonastre : « La voix n’est pas une biométrie classique »’. Laboratoire d’Innovation Numerique de la CNIL (LINC), 2 February 2017. https://linc.cnil.fr/fr/jean-francois-bonastre-la-voix-nest-pas-une-biometrie-classique.

Venkatadri, Giridhari, Elena Lucherini, Piotr Sapiezynski, and Alan Mislove. ‘Investigating Sources of PII Used in Facebook’s Targeted Advertising’. Proceedings on Privacy Enhancing Technologies, no. 1 (2019): 227–44. https://doi.org/10.2478/popets-2019-0013.

Von Schomberg, René. ‘The Precautionary Principle: Its Use Within Hard and Soft LawOR’,. European Journal of Risk Regulation 3, no. 2 (2012): 1447–156. www.jstor.org/stable/24323208.

Wachter, Sandra, and Brent Mittelstadt. ‘A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI’. Columbia Business Law Review Forthcoming (2019). https://ssrn.com/abstract=3248829.

Warren, Samuel, and Louis Brandeis. ‘The Right to Privacy’. Harvard Law Review 4, no. 5 (15 December 1890): 193–220. http://groups.csail.mit.edu/mac/classes/6.805/articles/privacy/Privacy_brand_warr2.html.

Weaver, Meaghann. ‘The Double Helix: Applying an Ethic of Care to the Duty to Warn Genetic Relatives of Genetic Information: The Duty to Warn Genetic Relatives of Genetic Information’. Bioethics 30, no. 3 (March 2016): 181–87. https://doi.org/10.1111/bioe.12176.

Weinberg, Alvin M. ‘Can Technology Replace Social Engineering’. Bulletin of the Atomic Scientists 22, no. 10 (1966): 4–8.

Weinreb, Lloyd L. ‘The Right to Privacy’. Social Philosophy and Policy 17, no. 2 (2000): 25–44. https://doi.org/10.1017/S0265052500002090.

Westin, A. F. Privacy and Freedom. New York: Atheneum Press, 1967.

Whitaker, Reginald. The End of Privacy: How Total Surveillance Is Becoming a Reality. New York: New Press : Distributed by W.W. Norton, 2000.

Whitley, Edgar A. ‘Informational Privacy, Consent and the “Control” of Personal Data’. Information Security Technical Report 14, no. 3 (August 2009): 154–59. https://doi.org/10.1016/j.istr.2009.10.001.

Whitman, James Q. ‘The Two Western Cultures of Privacy: Dignity versus Liberty’. Faculty Scholarship Series, no. 649 (2004). https://digitalcommons.law.yale.edu/fss_papers/649.

Witt, Charlotte. Ways of Being: Potentiality and Actuality in Aristotle’s Metaphysics. Ithaca, N.Y: Cornell University Press, 2003.

Xu, H., H. Wang, and A. Stavrou. ‘Privacy Risk Assessment on Online Photos’. In Research in Attacks, Intrusions, and Defenses 18th International Symposium, RAID 2015, Kyoto, Japan, November 2-4, 2015: Proceedings, by Herbert Bos, F. Monrose, and G. Blanc. Cham: Springer, 2015.