

Beyond the Concept of Anonymity: What is Really at Stake?

Björn Lundgren

The aim of this chapter is to discuss anonymity and the threats against it – in the form of de-anonymisation technologies. The question in the title above is approached by conceptual analysis: I ask what kind of concept we need and how it ought to be conceptualised given what is really at stake. By what is at stake I mean the values that are threatened by various de-anonymisation technologies. It will be argued that while previous conceptualisations of anonymity may be reasonable – given a standard lexical, or common-sense, understanding of the term – the concept of *anonymity* is not sufficient given what is really at stake. I will argue that what is at stake is *our ability to be anonymous*, which I will broadly characterise as *a reasonable control over what we communicate*.

The concept of anonymity has not – unlike the closely related concept of privacy – been given much attention by philosophers. The word ‘anonymous’ comes from the Greek *anōnumos*, which originally meant ‘without name’ or ‘nameless’.¹ However, as Kathleen A. Wallace notes, this is not the only sense of anonymity, neither is it the most interesting one (Wallace 1999: 23). According to Wallace a person P is anonymous, in a given respect, when others cannot coordinate traits so that P can be identified (ibid.: 25). Thus, we can conclude that on Wallace’s view anonymity is *non-identifiability* or, technically, *‘noncoordinatability of traits in a given respect(s)’* (ibid.: 25). While Wallace thinks of her definition as conceptual or metaphysical, her aim is also that it should enable us to analyse what is at stake (ibid.: 34).

Before Wallace, Helen Nissenbaum (1999) argued for a narrower conception of anonymity. According to Nissenbaum, what is at stake is

non-reachability, that is, what matters is not whether you can identify me, but whether you (after identification) can reach me. On the one hand, Nissenbaum's narrower conception makes sense if we think of, for example, criminals, who occasionally may care less that we know they commit a crime, as long as they can avoid punishment. On the other hand, it is also reasonable to argue that it is too narrow, given that the task of avoiding reachability may incur harm because one is identified.

Finally, Steve Matthews (2010), while discussing anonymity in relation to the social self, argues for another conception. According to Matthews, anonymity is non-trackability. What matters is neither only not being identified nor only not being reachable, but not being trackable over our social contexts. Matthews illustrates his point using examples discussed by Laud Humphrey about male-to-male encounters in public bathrooms. Matthews points out that these men were both partly identified and reachable and that what mattered to them was not avoiding identification or reachability in the context of their encounters, but that they could not be tracked from one social context to another. Matthews thinks of non-trackability both in the physical sense of not being tracked down and in the more abstract sense that people cannot link a piece of information to a more complete account of who that person is (2010: 354ff.), which corresponds to both Wallace's and Nissenbaum's conceptions of anonymity.

The aim in this chapter is not to question the above conceptualisations of anonymity, nor to argue which is superior. They all give a fair account of a common-sense conception of anonymity. What I will question is instead whether they give a fair account of what is really at stake. By what is at stake I mean the values that are threatened by various de-anonymisation technologies. In order to fulfil this aim, I will simplify the above accounts and focus on some fundamental features they all share.

First, on these views, anonymity is a bivalent concept (that is, a person either is or is not identifiable, reachable or trackable). In the next section, I will show that anonymisation technology is probabilistic. However, this alone does not imply that there is anything wrong with focusing on a conception of the state of anonymity. Indeed, previous accounts do recognise that de-anonymisation is probabilistic (for example, the non-coordinability of traits on Wallace's view). Thus, in the upcoming sections I will also show that the risk to anonymity matters, not only because it increases the likelihood of de-anonymisation (or author re-identification), but because it affects our ability to be anonymous and because it can cause harm beyond de-anonymisation.

Secondly, on these views, anonymity is a relation between some information (which I take to include actions, since actions can be expressed in the form of information) and a person, that is, what is at stake is revealing that some information is linked to some person (henceforth I will simply speak of 'revealing linkages' as 'linkages'). Conversely, I will argue that a person can be harmed when information about that person is linked to other information about that person, without it being linked to that person. We can distinguish these two linkages by referring to the former as 'information-to-person linkage' and the latter as 'information-to-information linkage'.

I will summarise the above theses in an account that I will call *the common view* (since it seems to correspond with the common-sense idea of anonymity):

The common view: a person is anonymous relative to some information only if that person cannot be singled out as *the one* linked to that information.

Before I turn to show that this is not the concept we need, two further things need to be noted.

First, the common view uses 'information' in a broader and more inclusive sense than what follows from a close reading of the three considered accounts. Secondly, although there are various senses of anonymity, it should be clear from the discussion above that the common view corresponds to the colloquial language notion of anonymity as relating to some *particular* information (that is, one is anonymous relative to some particular information, as compared with being a person about which very little is known).

The remainder of this chapter is structured as follows. In the next section, I will introduce some basic information about de-anonymisation technologies and an elementary challenge from Paul Ohm (2010). At the end of that section, I will have introduced reasons to doubt whether focusing on anonymity is enough. However, the choice of what concepts that we need will still be open. Thus, in the third section I argue that the common view cannot make sense of future risks of de-anonymisation and, more importantly, that harm goes beyond de-anonymisation, concluding that what we need is a concept of *our ability to be anonymous*. In the fourth section, I will aim to explicate this ability. Last, in the fifth section, while the previous discussions will have focused on individual harms, I aim to briefly broaden the discussion to show that the concept can be applied to collectives as well as to situations when de-anonymisation is all-things-considered desirable. Finally, I will end the chapter by a brief summation.

De-anonymisation and Information Aggregation

We can speak of a person being de-anonymised in line with the common view as follows:

The common view of de-anonymisation: a person that was previously anonymous relative to some information is de-anonymised if that person can be singled out as *the one* linked to that information.

In this sense de-anonymisation is achieved by linking information that allows for the identification of a singular person. For example, if you know the ZIP code, sex and date of birth of a someone living in the US, then there is at least 61 per cent chance that the person can be identified (see Ohm 2010: 1705, fn. 4, which also includes relevant references).

The aim of this section is to clarify a few relevant technical aspects of de-anonymisation technologies. First, I will discuss the granularity of the information that can be used to de-anonymise someone. Secondly, I will discuss the aggregating power of de-anonymisation technologies and set the basis for my critique of the common view.

In the previous section I noted that Wallace's conception of anonymity as *non-coordinatability of traits*, indicated a narrower scope than the common view. Wallace explains that traits should be understood as similar to that of a definitive description (Wallace 1999: 25). Furthermore, she thinks that 'a person is each of her traits (mother, novelist, English speaker, professor of mathematics, licensed driver, social security contributor, and so on)' (ibid.: 26).

Thus, conceptualising anonymity in terms of traits seems to imply that only information that is fairly substantial in itself (that is, information that is constitutive of who we are) can be used to de-anonymise someone. This is problematic since de-anonymisation can be performed using information of extremely varying granularity. It is misleading to think that only information that is constitutive of who we are can threaten our anonymity.

Indeed, de-anonymisation technology shows that information that is seemingly insignificant can be combined to de-anonymise us. For example, Johansson, Kaati and Shrestha (2015) have shown how de-anonymisation can be achieved by analysing *when* a user posts information. By analysing individual posting times, they create a 'timeprint' that can be used to reveal who the author is. With machine learning the success rate of author identification was 90 per cent for a sample of 1,000 users (ibid.: 10). They also show how to improve the identification by adding a stylometric analysis (that is, an analysis of how we write).

Although we can have a meaningful discussion about whether one's writing style constitutes a trait, it is clear that a singular time when one posts on the Internet does not constitute a trait in the sense Wallace conceptualises it (although a timeprint might). Despite this, it is the singular posting times that are coordinated into a timeprint, which allows for de-anonymisation.

Thus, we need to keep in mind that virtually *any* information, irrespective of its granularity, can be used to de-anonymise us. Therefore, to protect anonymity, we cannot only protect information that is sensitive in itself, since seemingly insignificant information can – if combined – become significant. This must be properly considered when conceptually analysing anonymity, which is why I have defined the common view in terms of information (and not in terms of, for example, traits).

Seemingly insignificant information can also be used to reveal further information about identified persons. An illustrative example is that of Facebook 'likes' (Facebook 'likes' is a function that allows users on Facebook to react to postings – including links – on Facebook by giving a thumbs-up; in 2016, Facebook complemented the function with a wider set of reactions). Kosinski, Stillwell and Graepel (2013) built a model that could analyse binary sorted personal traits on the basis of what people had 'liked' on Facebook. The success rate was between 60 and 95 per cent and included (in order of correlation – from weakest to strongest): parents together/separated before age of 21; drug use; single/in a relationship; alcohol use; smoker; lesbian; Christian/Muslim; Democrat/Republican; gay; gender (Kosinski et al.: 5803, fig. 2). They also tested correlation to various personal traits (ibid.: 5804, fig. 3).

This example illustrates that the kind of technologies that are used to de-anonymise some anonymous information about a person can also be used to reveal further information about that person, which is not part of any anonymous set of information. That is, what is at stake here is not merely retaining anonymity of some *particular* information, because these technologies can be used to generally reveal information about a person. The best way to conceptualise this harm is arguably to say that the problem is that the Kosinski et al. model in conjunction with the Facebook 'likes' affects a person's *ability to be anonymous*, because it diminishes a person's ability to control what is known about that person. However, this does not mean that the disvalue of having more information about oneself available *cannot* be explained on the common view. For example, one could view this as a de-anonymisation of a large set of information. However, on the one hand, the model neither identifies nor makes a person reachable (providing a potential counter-example to Wallace's and Nissenbaum's conceptions). But, on the other hand, it could be argued that it makes a person more

trackable over social contexts. The current question, then, is whether it is more reasonable to think of this as affecting our ability to be anonymous or our anonymity.

The above examples also relate to what Paul Ohm calls the ‘accretion problem’. The accretion problem is that successful de-anonymisation functions as a key to further success. By linking ‘anonymised’ information with other ‘anonymised’ information one can de-anonymise (or re-identify, which is the concept that Ohm uses) further ‘anonymised’ information. Generally, success increases the chance of further success (Ohm 2010: 1746). In a nutshell, the information gained by de-anonymising someone’s information in one situation can be further used to unlock other information to further de-anonymise someone’s information in other situations.

As Ohm notes, the accretion problem neatly illustrates why successful de-anonymisation of insignificant information matters: it matters because it can be used to de-anonymise significant information and cause future harm (Ohm 2010). It also illustrates why the risk of de-anonymisation is important. That an ethical analysis must take risks into consideration, rather than only focusing on idealised cases, has been pointed out before (see, for example, Hansson 2003).

As previously noted, on the common view anonymity is a binary concept, but that does not necessarily make de-anonymisation non-probabilistic.

On the one hand, it could be argued that the relevant concept is a measurement of the risk of de-anonymisation, not whether someone is anonymous or not. For example, Ohm argues that we should apply a concept of informational entropy (corresponding to some conceptions of information and informativity; see Adriaans 2013, for an overview). Entropy functions as a measurement of possibilities, fewer possibilities imply a lower entropy, which implies a greater risk of successful de-anonymisation. When more information is linked, the possible links (to a person or other information) become fewer and, therefore, de-anonymisation becomes more probable.

On the other hand, a proponent of the common view may respond that what matters is anonymity and that the risk of anonymity matters only insofar as it affects our anonymity. Thus, the risk of de-anonymisation is only a disvalue insofar as it is actually realised. Proponents of the common view can then use Ohm’s entropy-measurement (or something similar) to measure that risk.

However, harm does not stop with the *risk of de-anonymisation*. In the next section, I will argue that what is at stake here goes beyond the *risk of de-anonymisation*. First, we need a concept that also takes *future risks due to de-anonymisation* into consideration. Secondly, what is at stake goes beyond

de-anonymisation as such, because there can be harm from de-anonymisation technologies without de-anonymisation.

The Risk of Future De-anonymisation and Harm Without De-anonymisation

In the previous section, I relied on insights from Ohm – and others – that showed how insignificant linkages can be instrumental for the de-anonymisation of more sensitive information. In this section, I will first show that Ohm's suggestion to use entropy is insufficient given future threats to our ability to be anonymous. Next, I will more broadly attack the common view presumption that what is at stake is only information-to-person linkage.

As previously noted, Ohm suggested is that the risk of de-anonymisation can be measured by the concept of entropy:

... entropy measures how close an adversary is to connecting a given fact to a given individual. It describes the length of the interference chains heading in opposite directions, quantifying the remaining uncertainty. (Ohm 2010: 1749)

However, the problem is that while this gives us a measurement of the risk of de-anonymisation for a *given fact*, it cannot be used to measure future risks, given future *possible facts* (that is, something that would be a fact if a certain possibility is actualised in the future). Information-to-information linkage does not only affect the person's *current* ability to be anonymous, but also their *future* ability to be anonymous.

Here is an example: Jane has previously made a list of her favourite films available. While the list is clearly linked to Jane, there is no further anonymised information about Jane to which it could possibly be linked. Thus, on the common view, the potential harm of making her list known is limited to the harm of her list being known. On the view I defend, the existence of her list is also harmful because it affects her ability to be anonymous. The reason for this is simple: although it cannot currently be linked to any further information, it could be linked to future possible information (making Jane de-anonymised given her future actions). For example, consider that Jane creates an account on an online site that allows her to grade and rank movies (for example, the Internet Movie Database (IMDb)). Furthermore, suppose that her top-list (perhaps in conjunction with other previously available information) makes her IMDb account linkable to her. The example can be expanded, for example, by linkage from grading of all her films to making all her view-patterns available.

Suppose, for example, that she starts using an online streaming service such as Netflix. Further suppose that Netflix makes all their user data publicly available in an anonymised form (as they previously have, see Ohm 2010: 1720ff. for an over-view and a list of relevant references). If so, her currently available information would make it impossible for her, in the future, to keep her Netflix usage anonymous. As a result of that all of her viewing patterns would be available, which could be linked to reveal further information, etc. While de-anonymisation of her IMDb account may be an insignificant harm, in itself, her complete viewing-pattern (including both what she sees and when she sees it) would be possibly significant. Also, keep in mind that this further decreases her ability to be anonymous (via the accretion problem), both in the current situation as well as future possible situations.

On the common view, the information-to-person linkage between Jane and her IMDb account, via her top-list, would be the harm. On my view, that conclusion misses the fact that it was the previous release of her top-list that made the de-anonymisation and information-to-person linkage possible. On my view, we should reasonably consider *any linkage* harmful because it decreases our ability to be anonymous not only in the current situation, but in future possible situations. This is reasonably part of what is at stake. Conversely, a proponent of the common view would likely respond and argue that the harm here is future de-anonymisation of future possible facts, and that while we should recognise that information-to-information linkages was necessary for that harm, the information-to-information linkages was not harmful in itself, it was *only* instrumental for the de-anonymisation of information-to-person linkage. However, there can be harm from de-anonymisation technologies *without* de-anonymisation.

Consider the following example: suppose that a large set of sensitive information about Joe is linked so that it is knowable that it is about the same person, but not yet knowable whom it is about. Let us further presume that if Joe makes various kinds of communications, then these communications can be used to completely de-anonymise the sensitive information. Further, suppose that Joe knows this. This means that Joe would have to choose whether to communicate freely and, as a consequence, have a third party know that this information is about him, or whether to censor himself in order to keep it secret that the information is about him. (A similar example is discussed by Ohm 2010: 1748, however, Ohm's discussion is in line with the common view – that the de-anonymisation is the harm.)

In such a situation Joe is harmed, even if the information is never de-anonymised. Although this particular example may strike some as a rare

situation, it reveals that what matters goes beyond avoiding de-anonymisation as such. Furthermore, the example can be modified to generalise to situations most of us encounter in our daily lives. For example, choosing whether to create a Facebook account, use the services of Google, use some smartphone apps, or use many other services that come without a monetary cost, but include paying with your information. Many people do deal with such trade-offs. The problems are several, such as the non-transparent usage of (personal) information, which the example of Facebook 'likes' illustrates (see Hull, Lipford and Latulipe 2011 for a broad discussion, although partly outdated, about privacy concerns relating to Facebook). Furthermore, avoiding social networks such as Facebook may, at least in some social circles, include a certain amount of social isolation. While it is reasonable to think that many are unaware of the exact consequences for their personal information, most are not completely unwitting. There are plausibly a fair amount of people making trade-offs similar to that of Joe.

However, it is, of course, rare that a person knows of the *specific* risks of being de-anonymised. Thus, suppose, conversely, that Joe without knowing revealed all this sensitive information about himself (that is, he accidentally de-anonymised himself in relation to this information and, thus, revealed a large set of sensitive information about himself). While the common view can sensibly explain why the de-anonymisation matters in a situation in which Joe unwittingly reveals all that sensitive information, the core of the problem is not necessarily captured by the common view, but in control over our communications. Paraphrasing on Matthews' examples, suppose that the information is a detailed account of Joe's sexual activities. Suppose further that while the information may be sensitive (in the sense that it is not something that Joe wants to share), it does not reveal anything about his social position or role, which was not previously known. For example, Joe may have accidentally revealed details to friends about one of his sexual encounters (which he did not want to share), although it is perfectly well-known what *kind* of sexual encounters that Joe engages in. Thus, the information can be harmful, although it does not affect Joe's social position. Neither would it contribute to him being reachable in a sense that matters, which he previously was not. Nor would it make him identified (by revealing any of his traits), since he is already identified as person that does x-like things. Thus, the common view (or, specifically, the three discussed conceptions of anonymity) does not capture what is really at stake in these examples. In the next section, I will turn to explicate *our ability to be anonymous* and to show that it captures what is at stake.

Explicating Our Ability to be Anonymous

What the above examples show is that the values of anonymity go beyond the common view (that is, beyond information-to-person linkage). Harm can result simply from information-to-information linkage. Why? Because it decreases our ability to be anonymous. In this section, I aim to explicate the concept of our ability to be anonymous (in line with previous examples). I will also aim to show how it differs from the closely related concepts of privacy and the right to privacy.

The simple suggestion is that our ability to be anonymous should be defined in terms of *a reasonable control over what we communicate*. What I mean by control is actually absence of others' control over what we communicate, since control is not necessarily mutually exclusive (that is, that a person has control of x does not imply that other people do not also have control over x). Furthermore, by communicating I do not merely mean communication by speaking or writing. Rather, communication is much broader: what I have in mind is communicative actions, by which I mean actions that lead to information being transferred. For example, this can be speech acts, written messages, bodily movements (for example, how a person reacts in a situation or that a certain person goes from one place to another), or basically any usage of Internet services (just by clicking on websites we communicate, although such communication should in most cases ideally be private). Thus, what I broadly mean by a *control over what we communicate* is that when communicating some information, I, to a person, P, one has control to the extent that only I is communicated to only P. Relative to this, harm mainly relates to two main types of risk (which, as the previous examples show, need not be actualised – since risk can be part of a trade-off with another value): (1) harm can result because the information intended for P is also at risk of being spread to Q; and (2) harm can result from the risk of others' ability to conclude that I implies J. However, since most communications actually have such a deductive, inductive or abductive; the *reasonability* condition is quite central. If I ask you to pass me the dessert, you can perhaps presume that I like cake. This is, on most occasions, unproblematic. However, the currently available techniques make it possible to perform information linkage that goes well beyond such ordinary everyday conversational conclusions. The same holds for communications intended for P, but (at risk of being) spread to Q. This is, standardly, something that P has control over. But we can maintain a reasonable control if P respects that we should maintain this control.

If we re-examine the previous examples we can see that this ability indeed captures what was at stake. First, Jane's ability to control her future

communication was limited by unimportant information being linked to her. Because of information she was previously fine with communicating, it became impossible for her to keep other (more sensitive information) anonymous, which decreased her ability to be anonymous. Secondly, Joe had to choose between self-censorship and control over what was known about him (*mutatis mutandis* for the remaining variations), meaning that Joe was harmed, even without de-anonymisation, because his ability to be anonymous was seriously diminished.

Now, given that I talk about control over communications of (sensitive) information, it is easy to think that I am conflating our ability to be anonymous with privacy or the right to privacy. For example, according to Andrei Marmor, 'there is a general right to privacy grounded in people's interest in having a reasonable measure of control over the ways in which they can present themselves (and what is theirs) to others' (Marmor 2015: 3f.), and according to William A. Parent, '[p]rivacy is the condition of not having undocumented personal knowledge about one possessed by others' (Parent 1983: 269). However, if the overlap between these accounts of (the right to) privacy and our ability to be anonymous is, in fact, a conflation, then it is, arguably, these accounts – not mine – that are too broad. I have three things to say in defence of my conceptualisation: first, it is easy to conceive of an example that diminishes a person's control over what they communicate while having nothing to do with privacy. Suppose, for example, that Joe's sensitive information was his previously unknown business plans. Suppose further that although the business plans are sensitive they contain no privacy-sensitive information. Yet, by making it available, it affects his control over the ways in which he can present his business plan to others (that is, exactly the kind of control that Marmor discusses).² Thus, I would argue that it is control accounts of privacy (or the right to privacy) that are confused and in fact include too much (*mutatis mutandis* for Parent's account of possession of personal information). Secondly, the ability to be anonymous is reasonably also narrower than (the right to) privacy. Generally, there are situations which can diminish someone's privacy (or infringe upon the right) without affecting the control over our communications. For example, being seen in the nude by a former intimate partner standardly diminishes one's privacy (*mutatis mutandis* for the right), but given that it does not give the ex-partner access to any new information one's ability to be anonymous is not affected, since one's control over this information was already previously shared. Thirdly, that two concepts partly overlap is not necessarily a problem; it just means that our ability to be anonymous indeed has a lot to do with preserving privacy. Thus, while our ability to be anonymous is not

conflated with privacy, privacy is amongst one of many values that the ability can help protect (see Matthews 2010: 355).

That completes the main aim of the chapter, which was to argue for – and explicate – a conception of the concept we need in order to analyse and address possible harms from de-anonymisation technologies. However, the harms I have focused on so far have been limited to individuals' harm from de-anonymisation technologies. In the next section, I will aim to show that the concept of *our ability to be anonymous* applies more broadly, by addressing good and bad usage of de-anonymisation technologies.

Dual Usage of De-anonymisation Technologies

In this penultimate section, I will briefly contextualise the challenges we are facing, to show how the concept of our ability to be anonymous works as an umbrella concept that addresses a complex and multifaceted debate. While the examples previously discussed – in relation to the aim of chapter – have focus on situations when de-anonymisation technologies cause harm for individuals, these technologies are actually so-called dual-use technologies (that is, they can be used for good, bad and neutral purposes).

On the one hand, the ability to be anonymous broadly protects against any potential harm due to others' control of information about us. For individuals, this relates to fundamental questions about autonomy and liberty, because the ability not only protects privacy but also offers protection against undue influence and manipulation. Focusing on the latter, not previously discussed issues, there is a whole industry that depends on aggregation of individuals' online behaviour in order to adapt commercials to fit with individual preferences or to manipulate individual preferences to fit the product. What is worse is that if applied on an aggregated level, information aggregation can be used for gerrymandering, or targeted manipulations can be used to affect the political process (for example, by misleading people in order to make them vote for candidates they otherwise would not have supported). Although the dust is not yet settled on the 2016 US presidential election, some of the discussions concerning that election can serve as a good example of worst-case scenarios – when elections, indeed, are won by manipulating a large part of the electors. What I have in mind, amongst other things, is the scandal involving Cambridge Analytica, which 'harvested private information from the Facebook profiles of more than 50 million users without their permission' (Rosenberg, Confessore and Cadwalladr 2018).

On the other hand, the ability to be anonymous can be used for undesirable purposes. Indeed, sometimes de-anonymisation is desirable. For

example, we standardly want to de-anonymise criminals. De-anonymisation techniques can also be used to help determine who the author of an ancient text is (for example, whether *Alcibiades II* was written by Plato or not). Also, information aggregation of people's communications can also broadly be used to gain useful statistical information, such as of the spread of influenza (see, for example, Ginsberg et al. 2009). More importantly, the ability to be anonymous can be used to achieve a false identity, for example, to engage in disinformation campaigns. Here again, the US presidential election can be used as an example (see, for example, Mueller 2019). However, as previously implied, a successful disinformation campaign also, to some degree, depends on having sufficient knowledge about your target.

While the concept clearly applies to the above examples, one may worry that the concept is not analytically helpful when addressing trade-offs between good and bad uses of the ability. One way to view the challenges we are facing is as a binary choice between good and bad usage of the ability to be anonymous and de-anonymisation technologies (that is, similar to the debate between privacy and security). On the basis of this supposition, one cannot protect ordinary individuals' ability to be anonymous without protecting criminals' ability to be anonymous. Lastly, I will argue that this supposition is false.

Consider, for example, Onion Routing (used, for example, by the Tor browser). Onion Routing protects against a digital Peeping Tom by providing protection against an outsider's access to both the content of communications and the communicators (see, for example, Goldschlag, Reed and Syverson 1999). Techniques such as Onion Routing can be used to protect against eavesdroppers for both good purposes (for example, protecting a whistle-blower against an oppressive regime) and bad purposes (for example, protecting terrorists plotting an attack).

While Onion Routing can be used for both good and bad purposes, it cannot give any protection against the traces the ordinary Internet user leaves by their online behaviour, such as using services that require you to log-in, contributing to online information, or because services log meta-data (such as access time). This is because the distinct difference between aggregation of available information (making it possible to conclude that I implies J) and eavesdropping of private conversations.

Arguably the challenges to our ability to be anonymous are complex and multifaceted. But given that the activities of an ordinary person differ from those of a criminal, it is not unreasonable to think that it is possible to have a society that better protects some aspects of our ability to be anonymous, without leading to increased protection of criminal activity. The development of the Internet of Things – which will equip ordinary everyday

electronics with smart sensors – will put everyone’s ability to be anonymous at risk. However, while it is not logically impossible that, for example, a smart fridge could collect information that is helpful in the solution of a crime, it is not very likely. Thus, although there are arguably trade-offs to each option (for example, there are obvious individual benefits of aggregated analysis of a person’s ‘communications’), the trade-offs are not necessarily between bad and good uses of people’s ability to be anonymous. Reasonably, successful criminals know better than ordinary citizens how to protect their ability to be anonymous (making use of techniques that few ordinary people apply). This does not only put ordinary individuals at a disadvantage, but – as the example with the fridge aims to illustrate – we can protect ordinary people’s ability to be anonymous in ordinary situations without necessarily protecting the ability to be anonymous for more nefarious purposes.

Either way, the account of the concept of our ability to be anonymous, rather than anonymity, is indeed the most helpful concept to analytically disentangle and evaluate these issues.

Summation and Final Comments

In this text, I have defended two intertwined ideas. First, given the power of the available de-anonymisation techniques the challenges we are facing are best conceptualised by an *ability to be anonymous*, rather than by a concept of *anonymity*. Secondly, our *ability to be anonymous* can be conceptualised in terms of having *reasonable control over what we communicate*.

De-anonymisation and information aggregation can be used to analyse our communications and behaviour in a way that allows others to infer more information than what is intended or available analytically in the communication as such. On the one hand, this can infringe upon our (right to) privacy, force us into lose-to-lose trade-offs and be abused to manipulate us, something which puts the whole democratic society at risk. On the other hand, it can be used to de-anonymise criminals, to gain valuable aggregated information (such as the spread of influenza), and to uncover the source of disinformation campaigns.

While the techniques are both beneficial and disadvantageous, it is reasonable to presume that it is possible to protect individuals against harm, while retaining various benefits.

Last, I have avoided saying what we should do, but given the current technological development it is not unreasonable to think that we are at a crossroads. We either choose to act in way that protects our ability to be

anonymous or we might move towards a society where the private sphere and our individual control of it, and possibly – as an extension – our autonomy, will be limited. Either way, what we need to analyse is not the concept of *anonymity*, but the concept of *our ability to be anonymous*.

Notes

1. See, for example, at: <https://www.lexico.com/en/definition/anonymous>.
2. It is important to point out that Marmor actually talks of the right to privacy as being *grounded* in an interest in a certain form of control, rather than being identical to that certain form of control. Thus, it is possible that Marmor would define the right to privacy in a different way. Hence, what is said above should not – without further qualification – be read as a criticism of Marmor's paper.

References

- Adriaans, P., 'Information', *The Stanford Encyclopedia of Philosophy*, Fall 2013 edn.
- Ginsberg, J., M. H. Mohebbi, R. S. Patel, L. Brammer, M. S. Smolinski and L. Brilliant, 'Detecting Influenza Epidemics Using Search Engine Query Data', *Nature*, 457 (2009): 1012–14.
- Goldschlag, D., M. Reed and P. Syverson, 'Onion Routing for Anonymous and Private Internet Connections', *Communication of the ACM* 42(2) (1999): 39–41.
- Hansson, S. O., 'Ethical Criteria of Risk Acceptance', *Erkenntnis* 59(3) (2003): 291–309.
- Hull G., H. R. Lipford and C. Latulipe, 'Contextual Gaps: Privacy Issues on Facebook', *Ethics & Information Technology* 13 (2011): 289–302.
- Johansson, F., L. Kaati and A. Shrestha, 'Timeprints for Identifying Social Media Users with Multiple Aliases', *Security Informatics* 4(7) (2015).
- Kosinski, M., D. Stillwell and T. Graepel, 'Private Traits and Attributes are Predictable from Digital Records of Human Behavior', *PNAS* 110(15) (2013): 5802–5.
- Matthews, S., 'Anonymity and the Social Self', *American Philosophical Quarterly* 47(4) (2010): 351–63.
- Marmor, A., 'What Is the Right to Privacy?' *Philosophy & Public Affairs* 43 (2015): 3–26.
- Mueller, R. S., 'Report on the Investigation into Russian Interference in the 2016 Presidential Election. Volume I', *US Department of Justice*, 2019, available at: https://upload.wikimedia.org/wikipedia/commons/e/e4/Report_On_The_Investigation_Into_Russian_Interference_In_The_2016_Presidential_Election.pdf, last accessed 28 September 2019.
- Nissenbaum, H., 'The Meaning of Anonymity in an Information Age', *The Information Society* 15(2) (1999): 141–4.
- Ohm, P., 'Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization', *UCLA Law Review* 57 (2010): 1701–77.

- Parent, W. A., 'Privacy, Morality, and the Law', *Philosophy & Public Affairs* 12 (1983): 269–88.
- Rosenberg, M., N. Confessore and C. Cadwalladr, 'How Trump Consultants Exploited the Facebook Data of Millions', *New York Times*, 17 March 2018, available at: <https://www.nytimes.com/2018/03/17/us/politics/cambridge-analytica-trump-campaign.html>, last accessed 28 September 2019.
- Wallace, K. A., 'Anonymity', *Ethics and Information Technology* 1 (1999): 23–35.