

4

Are Publicly Available (Personal) Data ‘up for grabs’? A Discussion of Three Privacy Arguments

ELISA ORRÙ*

Abstract

The re-use of publicly available (personal) data for originally unanticipated purposes has become common practice. Without such secondary uses, the development of many AI systems like large language models (LLMs) and ChatGPT would not even have been possible. This chapter addresses the ethical implications of such secondary processing, with a particular focus on data protection and privacy issues. Legal and ethical evaluations of secondary processing of publicly available personal data diverge considerably both among scholars and the general public. While some of these uses are met with opposition and criticism, others are quite unanimously viewed as unproblematic. Often, proponents and opponents of such practices invoke the same ethical and legal standards for their opposite conclusions. This state of affairs shows that other considerations besides the public availability of data must play a role. It calls for a theoretical clarification of the additional criteria that should guide decisions about the (legally informed) ethical acceptability of re-processing practices. In order to make a contribution towards this goal, the present chapter maps the ongoing debate and systematises the existing contributions around three lines of argument: a consent-centred position, an approach that focuses on the distinction between data and information, and finally a line of argument that focuses on the contextual norms that govern the flows of information. The chapter further relates these arguments to three underlying conceptions of privacy and data

*The research conducted for this chapter was partly funded by the Horizon Europe project VIGILANT (Vital IntelliGence to Investigate ILlegAl DisiNformaTion, Grant agreement ID: 101073921). I would like to thank Deborah Krzyzowski for her assistance with literature search, Brendan Spillane and two anonymous reviewers for their stimulating comments on earlier versions of the chapter. The chapter was presented and discussed at the Max Planck Institute for the Study of Crime, Security and Law in Freiburg. I am grateful to my colleagues at the Institute for their insightful comments.

protection: rights-based, structural and contextual, and discusses the advantages and disadvantages of each position in the light of concrete examples. It concludes by arguing for a mixed approach that combines core elements of the structural and contextual approaches. The chapter aims to contribute to existing research in the fields of data, AI and research ethics, and to reconnect the debate with ethical and legal scholarship on privacy and data protection. In doing so, it aims to make a theoretical contribution towards refining existing conceptions of privacy and data protection in order to make them more fit to ‘drive our digital world’ as far as the use of publicly available data is concerned.

Keywords

Privacy, Data protection, Contextual integrity, Large language models (LLMs), Web harvesting, Data mining, Social media research, Research ethics, Data ethics.

I. Introduction

Publicly available data are data that have been made available to the general public, usually by being published on the Internet. Such data may include information that is not personally identifiable or that is part of the public domain (such as Wikipedia articles), or they may include personal information, as is often the case with social media posts. The personal information shared may be made public by the individuals themselves or by others, as is the case, for example, when an individual posts pictures of several people.

These data are valuable resources for a wide variety of uses and purposes, ranging from advertising to research to the development of apps and other technologies for commercial exploitation. Therefore, they are often re-processed for purposes that were not foreseen at the time they were made public.

Commonly used practices of re-processing include the use of social media data to make assumptions about users’ mental health, the scraping of the web for individual portraits to build biometric datasets, and the use of written text to train LLMs. One example of such usage are apps that are specifically designed to detect early signs of depression or suicidal tendencies in individuals. For instance, the Samaritans’ Radar app allowed Twitter users to install a plug-in which would alert them if any of the profiles they had chosen to monitor showed signs of suicidal intent.¹ Other examples are studies aimed at identifying language patterns in suicidal people before and after suicide attempts, and predicting the likelihood of mothers developing postpartum depression by analysing their tweets.² Another

¹ ‘Samaritans Radar’, Samaritans, www.samaritans.org/about-samaritans/research-policy/internet-suicide/samaritans-radar/, accessed 12 October 2023.

² Glen Coppersmith et al., ‘Exploratory Analysis of Social Media Prior to a Suicide Attempt’, in *Proceedings of the Third Workshop on Computational Linguistics and Clinical Psychology* (San Diego,

type of secondary use of publicly available data concerns images of individuals, either on social media or on the websites of public organisations, private companies, etc. In 2020, the AI company PimEyes, originally based in Poland but now located outside the EU, and the US startup Clearview AI used facial images available on the internet to create their own biometric databases. They then provided a biometric search service that allows users to upload an image and search for matches (that is, images of the same person). This service is available for either private or law enforcement use.³ Concerning the third type of use, virtually all textual content available on the Internet, including books and scientific articles, Wikipedia entries, contents from blogs, chats, social media and websites, can be used to train LLMs. These models form the essential functions of multiple applications, including voice assistants, automatic translations and chatbots like ChatGPT.

While the Samaritans' Radar was discontinued a few days after its launch due to widespread criticism, other similar applications, such as the Live for Tomorrow chat service in New Zealand, have reportedly been widely accepted.⁴ Alongside enthusiastic acceptance, systems such as voice assistants and ChatGPT have been met with public criticism, although this is not primarily focused on the re-use of publicly available data as training data.⁵ Finally, the re-use of images for biometric identification, as practiced by PimEyes and ClerviewAI, has generated intense discomfort, especially in the EU, and triggered a series of coordinated complaints to data protection authorities in several EU countries, as well as the current attempt by the EU Parliament to add this type of activity to the list of prohibited practices under Article 5 of the draft AI law.⁶

CA, US: Association for Computational Linguistics, 2016): 106–17, doi.org/10.18653/v1/W16-0311; Munmun De Choudhury, Scott Counts, and Eric Horvitz, 'Predicting Postpartum Changes in Emotion and Behavior via Social Media,' in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13 (New York, NY, US: Association for Computing Machinery, 2013): 3267–76, doi.org/10.1145/2470654.2466447.

³Since March 2022, also the Ukrainian Ministry of Defence has been using Clearview AI's facial recognition technology for military aims: 'Exclusive: Ukraine has started using Clearview AI's facial recognition during war,' Reuters, published 14 March 2022, www.reuters.com/technology/exclusive-ukraine-has-started-using-clearview-ais-facial-recognition-during-war-2022-03-13/.

⁴Jennifer Nicholas, Sandersan Onie, and Mark E Larsen, 'Ethics and Privacy in Social Media Research for Mental Health,' *Current Psychiatry Reports* 22, no. 12 (2020): 1–7, doi.org/10.1007/s11920-020-01205-9, 2. It seems, however, that the original proactive functioning of the app as described by Nicholas et al. is no longer active and that Live For Tomorrow has mutated to a service for connecting users of social media platforms with mental health and crisis helplines. See 'Working to improve the mental health of people in New Zealand and around the world,' LifeForTomorrow, livafortomorrow.co/, accessed 12 October 2023.

⁵Ethical issues that have been discussed in reference to voice assistants regard for instance gender biases, while concerns on ChatGPT include issues of intellectual property, good research practices and the impact on educational systems. See for instance 'I'd Blush if I Could: Closing Gender Divides in Digital Skills through Education,' UNESCO and EQUALS Skills Coalition, 2019, unesdoc.unesco.org/ark:/48223/pf0000367416.page=1; 'CRITICAL AI: Adapting College Writing for the Age of Large Language Models Such as ChatGPT: Some Next Steps for Educators,' Critical AI, 17 January 2023, criticalai.org/2023/01/17/critical-ai-adapting-college-writing-for-the-age-of-large-language-models-such-as-chatgpt-some-next-steps-for-educators/.

⁶See 'Europaweite Beschwerden gegen Clearview AI,' noyb, 26 May 2023, noyb.eu/de/europaweite-beschwerden-gegen-clearview-ai and 'Parliament's negotiating position on the artificial intelligence act,' European Parliament, June 2023, www.europarl.europa.eu/RegData/etudes/ATAG/2023/747926/EPRS_ATA(2023)747926_EN.pdf.

A glance at the academic literature on the permissibility of such practices reveals an equally diverse set of positions. Legally, such re-uses are often justified by virtue of their being publicly available. Article 9(2)(e) General Data Protection Regulation ('GDPR') seems to support this interpretation, at least in some cases, as it even allows the processing of special categories of data if they are 'manifestly made public by the data subject'. However, some legal scholars interpret the same provision as not justifying secondary processing if the latter is aimed at extracting information that the data subject would not reasonably expect to share while making the data public.⁷

Similarly, from an ethical point of view, there does not seem to be a straightforward interpretation of the issues at stake. Despite the growing number of guidelines on data ethics and AI ethics, the issue remains an ethically grey area.⁸ Drawing on good scientific practice and ethical guidelines for research involving human participants, some scholars argue that the use of these data should be subject to additional guarantees including informed consent, while others argue that secondary use of personal data without consent is ethically permissible.⁹ Notably, ethical considerations regarding the use of such data are often not discussed at all, not only in the commercial sector, but also in research, where ethical self-assessment and oversight are increasingly important. A recent study analysed 132 research articles using data from social media and blogs for discourse analysis and found that two thirds of them did not report or discuss ethical issues whatsoever.¹⁰ When justifications were given, the public availability

⁷ Stephan Schindler and Gerrit Hornung, 'Datenschutz bei der biometrischen Gesichtserkennung,' *Datenschutz und Datensicherheit – DuD* 45, no. 8 (2021): 515–21, doi.org/10.1007/s11623-021-1482-6; Gerrit Hornung and Carolin Gilga, 'Einmal öffentlich – Für immer schutzlos?,' *Computer und Recht* 36, no. 6 (June 2020): 367–79, doi.org/10.9785/cr-2020-360609. See section III below for further details. Note that the so-called Digital Services Act is not going to change significantly the regulatory landscape in this respect, since the obligation it puts upon social media platforms to enable researchers to access to real-time publicly available data only for the purpose of 'performing research that contributes to the detection, identification and understanding of systemic risks in the Union pursuant to Article 34(1), Digital Services Act, Art 40(4). This means that the research to be conducted is not on the publicly accessible data themselves, but on the functioning of the services provided.

⁸ For an overview of existing ethics guidelines see Anna Jobin, Marcello Ienca, and Effy Vayena, 'The Global Landscape of AI Ethics Guidelines,' *Nature Machine Intelligence* 1, no. 9 (September 2019): 389–99, doi.org/10.1038/s42256-019-0088-2.

⁹ As an example of the former position see Signe Ravn, Ashley Barnwell, and Barbara Barbosa Neves, 'What Is 'Publicly Available Data'? Exploring Blurred Public–Private Boundaries and Ethical Practices Through a Case Study on Instagram,' *Journal of Empirical Research on Human Research Ethics* 15, no. 1–2 (February 2020): 40–45, doi.org/10.1177/1556264619850736, as an example of the latter position see Guusje Jol and Wyke Stommel, 'Ethical Considerations of Secondary Data Use: What about Informed Consent?,' *Dutch Journal of Applied Linguistics* 5, no. 2 (January 2016): 180–95, doi.org/10.1075/dujal.5.2.06jol.

¹⁰ However, this does not necessarily mean that ethical issues were completely ignored in the research. For example, the authors of the study report that, in most cases, the names of the authors of quoted posts were deleted. Although this practice does not reliably protect the users' anonymity (for instance in case of post retrieval), it shows that *some* ethical considerations played a role even if not explicitly discussed. See Wyke Stommel and Lynn de Rijk, 'Ethical Approval: None Sought. How Discourse Analysts Report Ethical Issues around Publicly Available Online Data,' *Research Ethics* 17, no. 3 (July 2021): 275–97, doi.org/10.1177/1747016120988767.

of the data was mostly seen as an appropriate and sufficient justification for further processing.¹¹

The divergence in public perceptions and expert opinions suggests that while the fact that the data are publicly available may play a role in determining the lawfulness and ethical acceptability of further processing, additional factors and considerations should be taken into account. This chapter aims to clarify which additional elements are crucial to distinguish cases where the further use of publicly available personal data is permissible from cases where it is not. To this end, it systematises the existing (still sparse) ethical and legal literature on the topic around three lines of argument that correspond to three underlying conceptions of privacy and data protection.

The first line of argument focuses on consent, the second on the distinction between data and information, and the third on informational contexts. According to the first line of argument, which can be seen as an example of rights-based approaches to privacy and data protection, a new right not to be profiled by AI that reuses publicly available data should be introduced, which can, however, be waived if data subjects consent to such processing.¹² The second kind of reasoning is based on the distinction between the data that is made public and the information that can be 'extracted' from it, and stresses the importance of supporting individual data protection rights with structural interventions to appropriately regulate and shape the way data are handled.¹³ Finally, the third line of argument focuses on the contexts in which information is exchanged and, drawing on conceptualisations of privacy as contextual integrity, emphasises the need to process data in accordance with the norms that govern information flows in the relevant contexts.¹⁴

¹¹ Stommel and Rijk, 'Ethical Approval: None Sought'. See also Joanna Taylor and Claudia Pagliari, 'Mining Social Media Data: How Are Research Sponsors and Researchers Addressing the Ethical Challenges?', *Research Ethics*, 26 October 2017, doi.org/10.1177/1747016117738559.

¹² For a characterisation of rights-based approaches see Daniel J Solove, 'The Limitations of Privacy Rights', *Notre Dame Law Review* 98 (2023): 975–1036, dx.doi.org/10.2139/ssrn.4024790. The introduction of the right not to be profiled through AI has been suggested by Thomas Ploug, see 'The Right Not to Be Subjected to AI Profiling Based on Publicly Available Data—Privacy and the Exceptionalism of AI Profiling', *Philosophy and Technology* 36, no. 1 (2023): 1–22, doi.org/10.1007/s13347-023-00616-9.

¹³ As examples of this line of argument see Nicholas, Onie, and Larsen, 'Ethics and Privacy in Social Media Research for Mental Health'; Schindler and Hornung, 'Datenschutz bei der biometrischen Gesichtserkennung' and, for the underlying conception of data protection, Marion Albers, 'Umgang mit personenbezogenen Informationen und Daten', in *Grundlagen des Verwaltungsrechts*, eds. Andreas Voßkuhle et al., vol. 2 (München: C.H. Beck, 2006), § 22; Marion Albers, *Informationselle Selbstbestimmung* (Baden-Baden: Nomos, 2005), doi.org/10.5771/9783845258638; Marion Albers, 'Realizing the Complexity of Data Protection', in *Reloading Data Protection*, eds. Serge Gutwirth, Ronald Leenes, and Paul de Hert (Dordrecht: Springer, 2014), 213–35, doi.org/10.1007/978-94-007-7540-4_11.

¹⁴ See Hannah Brown et al., 'What Does it Mean for a Language Model to Preserve Privacy?', *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, FAccT '22 (New York, NY, USA: Association for Computing Machinery, 2022), 2280–92, doi.org/10.1145/3531146.3534642 and, for the theoretical background, Helen Nissenbaum, *Privacy in Context: Technology, Policy, and the Integrity of Social Life* (Stanford: Stanford Law Books, 2010); Helen Nissenbaum, 'Respecting Context to Protect Privacy: Why Meaning Matters', *Science and Engineering Ethics* 24, (2018): 831–852, doi.org/10.1007/s11948-015-9674-9.

The present chapter discusses the pros and cons of each position in relation to concrete examples from the three areas mentioned above (sections II-IV), and proposes a fourth, mixed approach that combines key elements of the second and third types of conceptualisation (section V).

II. An Individual Rights Approach: Consent Required

The first line of argument relies on consent, and considers the re-processing of publicly available data to be legitimate only if data subjects have consented to this secondary use. An argument along these lines has been formulated by Thomas Ploug in relation to re-processing for AI profiling purposes.¹⁵

Ploug argues for the introduction of a *sui generis* (ie, peculiar) right not to be subjected to AI profiling based on data that have been made public, for example on social media platforms. The need for such a distinctive right prohibiting AI profiling derives from the specific harms that can be caused by AI-based profiling, such as undue exposure to social control, stigmatisation and self-stigmatisation.¹⁶ The author discusses cases in which profiling is used to infer information about the mental health of data subjects, but argues for a legal right not to be profiled by AI in general, regardless of the specific profiling purpose.

This right is conceptualised as a negative, *pro-tanto* (ie, non absolute) right. As a negative right, it is characterised as a right to non-interference through AI profiling. As a *pro-tanto* right, it allows for exceptions, limitations and balancing against other rights. Furthermore, this right can be waived by data subjects if they consent to AI-based profiling.¹⁷ Ploug's suggestion presents structural similarities with some GDPR provisions such as Article 9(2)(a) GDPR, which allows for exceptions to the general prohibition of processing sensitive data if data subjects explicitly consent to the processing.¹⁸ However, according to Ploug, the GDPR only protects individuals from AI profiling indirectly, by prohibiting certain types of profiling, but does not specifically entail a right of the kind he proposes.¹⁹

Ploug's position can be seen as prototypical of rights-based conceptualisations of privacy. These understandings rely on individual rights as the primary means of ensuring privacy. However, as Daniel Solove has noted, these approaches suffer from a number of limitations.²⁰ In a sense, privacy rights

¹⁵ Ploug, 'The Right Not to Be Subjected to AI Profiling.'

¹⁶ *ibid* 7–11.

¹⁷ *ibid* 2.

¹⁸ 'the data subject has given explicit consent to the processing of those personal data for one or more specified purposes.'

¹⁹ For more details see Ploug, 'The Right Not to Be Subjected to AI Profiling,' 15–17.

²⁰ Solove, Daniel J., 'The Limitations of Privacy Rights.' For an overview of other common critiques to individual approaches to privacy and especially on notice-and-consent frameworks see

demand too much of individuals: first, because they rely on individuals to address problems that are systemic, and second, because individuals often lack the time and knowledge to make effective use of their rights. Finally, privacy cannot be protected by focusing only on individuals in isolation: sharing one's data may also mean sharing information about other people (for example, by sharing their own genomic information, individuals also share information about their relatives).²¹

Ploug's proposal mirrors these general difficulties of individual rights conceptualisations in three ways. The first limitation of his approach resides in the workability of informed consent: how can individuals meaningfully exercise this right if data are already 'out there'? How can they check whether their data are being used for profiling by AI, possibly years after they have made them public, and by actors and purposes of which they are not even aware? How can processors, who want to further analyse the data, know whether the data subjects gave their consent for such re-use when they made their data public?

A second limitation, that Ploug himself acknowledges, is that making one's data public and using them for profiling may also allow inferences to be made about others, such as relatives who may have similar predispositions to diseases. However, the author does not clarify how this can be mitigated by a consent approach if consent is not given by people indirectly affected by the processing activities.

Finally, on a general level, Ploug's proposal lacks systematicity because it formulates a very specific right that applies only to cases where the personal data is used for profiling by AI, and does not include a general caution against other kinds of processing. Even if (as I believe) not all uses of publicly available personal data are problematic, it might be reasonable not to limit the focus from the outset to one particular kind of processing. Given that the ways in which data can be reused (through AI applications and other means) are constantly evolving and, for the most part, unpredictable, it seems more advantageous to look first at the processing of publicly available personal data as such and, in a second step, to investigate which criteria make it possible to identify, among all the possible uses, those that are not ethically problematic. Otherwise, there is a risk of ending up with a reactive, patchwork set of standards that will not be able to cope with new challenges and developments.

Effective protection against the risks of re-processing of personal data can therefore, in my view, only be achieved by moving away from a focus on the individual and towards a more structural approach which emphasises the role of institutions. This would aim at regulating the architecture of data exchange in a

also Daniel Susser, 'Notice After Notice-and-Consent: Why Privacy Disclosures Are Valuable Even if Consent Frameworks Aren't,' *Journal of Information Policy* 9 (2019): 148–73, doi.org/10.5325/jinfopoli.9.2019.0148.

²¹ Solove, 'The Limitations of Privacy Rights,' 975.

way that creates the conditions for data to be handled appropriately, without relying too much on the initiative of individuals to enforce their rights. Additionally, a systematic approach is preferable to one that creates ad hoc individual rights focusing on specific processing purposes such as profiling by AI.

III. A Structural Approach: Data and Information are not the Same

The second line of argument takes a structural approach and distinguishes between the *data* that have been made public and the *information* that can be extracted from them. This information may reveal further characteristics, attitudes, etc about individuals that were not intended to be shared or communicated when the data was originally made public.

Recently, Jennifer Nicholas et al.²² have adopted a similar line of argument to discuss the use of social media data to infer information about users' mental health. They observe that users make statements on social media that are directly or indirectly related to mental health issues and that are analysed to draw inferences about mental illness contrary to their original intentions.²³ When sharing the original data, social media users want to convey information about, say, their feelings and emotions, but most often do not want to provide information to researchers or health care workers about their mental health or illnesses. The authors therefore suggest that re-use should be limited to data that have been explicitly donated by social media users.

Another illustrative example is the use of publicly available portraits to extract biometric information, as discussed by Stephan Schindler and Gerrit Hornung.²⁴

The authors distinguish between the data that internet and social media users intended to make public and the information that can be extracted from such data. By posting their portraits, the individuals did not intend to share the biometric information that is technically embedded in their images.²⁵ Crucially, according to Schindler and Hornung, the very fact that a special technical process is required to extract the biometric information from the portraits provides

²² Nicholas, Onie, and Larsen, 'Ethics and Privacy in Social Media Research for Mental Health.'

²³ '[...] although social media data are often publicly available, when collected for the purposes of mental health research, the *data* are likely to contain sensitive personal *information*. For example, as outlined above, social media data may detail an individual's experience with a mental health condition, or be used to infer mental ill-health when no such public (or potentially private) declaration has been made,' Nicholas, Onie, and Larsen, 'Ethics and Privacy in Social Media Research for Mental Health,' 3, italics added.

²⁴ Schindler and Hornung, 'Datenschutz bei der biometrischen Gesichtserkennung.'

²⁵ The authors discuss here only the cases in which the portraits are made publicly accessible by the same subjects depicted on them. In all other cases there is no legitimate ground for processing the data, see Art 6 and Art 9 GDPR.

evidence against the data subjects' willingness to make this information public by sharing their images.²⁶

Interestingly, the authors ground their argument in an interpretation of Article 9 GDPR and thus see existing positive EU law as capable of protecting individuals from such uses (whereas, as we have seen above, Ploug considers the existing legal protection to be insufficient and argues for the introduction of a *new* legal right). Biometric data, according to Article 9 GDPR, are among the special categories of personal data whose processing is generally prohibited. Exceptions to this general prohibition are possible if, in addition to a valid legal basis under Article 6 GDPR, one of the special conditions listed in Article 9(2) GDPR applies. The only possible exception that could in principle apply to the case study concerns cases in which 'processing relates to personal data which are manifestly made public by the data subject' (Article 9(2)(e) GDPR). However, as we have seen, Schindler and Hornung exclude that the biometric information contained in the images was 'manifestly made public' by the data subject, who most likely did not consider the possibility of extracting this information from their pictures (also because, in many cases, this extraction was not technically possible at the time the images were published on the Internet).

Even if Schindler and Hornung do not consistently distinguish between 'data' and 'information' terminologically, they do so conceptually. The distinction between data and information is not new in German legal thought, especially in discourses on the right to informational self-determination, and has been elaborately conceptualised by Marion Albers.²⁷ According to Albers, data are the 'underpinnings' or bases for information, while 'information can only be understood in context'.²⁸ More specifically, 'data are signs that are recorded on a data carrier and can function as information bases', while 'information is constituted by meaningful items that are generated in a particular social context from observations, communications or data'.²⁹ Data therefore convey information, but information is more than data. Data become information only in a specific context and when interpreted by actors who make sense of them.

It is important to note that, according to Albers, data protection law usually does not recognise this distinction. In fact, the failure to make such a distinction is one of the main reasons for the lack of effective protection of individuals against improper use of information about them. Such a distinction would make it possible to recognise that the actual object of data protection law is not data as such, but persons. Thus, the critical question that allows a particular need for protection to

²⁶ Schindler and Hornung, 'Datenschutz bei der biometrischen Gesichtserkennung', 521.

²⁷ Albers, 'Umgang mit personenbezogenen Informationen und Daten' and Albers, *Informationelle Selbstbestimmung*.

²⁸ Albers, 'Umgang mit personenbezogenen Informationen und Daten', 113. Translation by the author, assisted by AI.

²⁹ Albers, 'Umgang mit personenbezogenen Informationen und Daten', 115, translation by the author, assisted by AI.

be identified is not what type of data are used, but rather what information about individuals can be derived from those data. The latter question, in turn, can only be answered by reference to the contexts in which the meaning of the information unfolds.³⁰

Crucially, the right to data protection, reformulated as a right to adequate treatment of personal data and information, is not to be understood as a classical negative right. It does not only require the state to refrain from interfering with the individual's sphere of liberty. On the contrary, it also obliges the state to actively regulate the flow of information and to shape it in such a way that individuals can effectively make use of the relevant fundamental rights.³¹

The example of secondary processing of social media data to build hypotheses about users' mental health can now be framed more precisely as follows: in making their statements on social media platforms, users share data (the individual characters typed and recorded on social media servers) to convey information such as 'I feel happy today' or 'I was disappointed by my friend's behaviour'. However, they do not intend to convey information such as 'I'm probably not depressed' or 'I'm paranoid', nor do they intend to allow other people to infer this information from their posts.

Similarly, in relation to the uses of portraits discussed by Schindler and Hornung, we can now frame the issue more precisely: when individuals post their images in freely accessible venues on the Internet, they share data (the pixel of the image they upload) with the intention of conveying specific information calibrated to the given context. For example, if they share a picture of themselves at a party, the information to be conveyed may be: 'last evening I was there' or 'look what an interesting life I have!'; If they upload an image on their professional website as a freelancer, they probably want to convey the information: 'this is how I look like (intelligent, friendly, reliable ...): I can be trusted' and so on. In all these cases, they most probably did not intend to share any information about the biometric features of their faces, even if such information can be extracted from the pixel contained in the image they published.

The distinction between data and information therefore seems useful in clarifying the issue at stake. The core of the matter is that instead of focusing on the qualities of the data shared by users, we should concentrate on the characteristics of the information that can be extracted from them: what does this information reveal about a particular individual?

Moreover, this approach allows us to sidestep the difficult question of the consent of data subjects and to focus on the consequences of data processing. It allows also to adequately conceptualise cases where sensitive information is derived from non-sensitive data or information acting as proxies: for example,

³⁰ Albers, 'Umgang mit personenbezogenen Informationen und Daten,' 124–25.

³¹ Albers, *Informationelle Selbstbestimmung*, 605.

when 'ethnicity' is inferred from a postcode (a non-sensitive information).³² The distinction between data and information takes into account the fact that data carry more information than that which is intended for use in the contexts in question: if data have been shared ('made public') for the purpose of conveying a particular piece (or pieces) of information, but it is used to obtain other, sensitive, information, a breach of privacy occurs.

Within this approach, however, a number of questions remain unanswered: does the right to adequate treatment of personal data, as formulated by Albers, require additional protection when the information that can be derived from the data is 'sensitive' in the sense that it relates to race, sex, political or religious opinions etc (thus broadly corresponding to the special categories of data under Article 9 GDPR), or should such protection apply to any type of information about individuals that was not originally intended to be shared?

Schindler and Hornung's specific argument seems to apply only to sensitive information. Moreover, they ground their argument on the necessity of an additional technical procedure to 'extract' biometric information from the original data. Is the use of additional technical procedures then a necessary condition for distinguishing between legitimate and illegitimate secondary processing and information extraction? And if so, are all kinds of data processing to be understood as additional technical procedures in this sense, or only certain kinds of data processing? And if the latter, does big data analysis fall into this category, or only machine learning methods or any use of AI?

Finally, the authors rightly assume that the sensitivity of an item of information can only be determined by reference to the context of its use and is not a fixed characteristic of the data itself. However, this seems to imply that the special protection required for the processing of special categories of data can only be determined *ex post* and on a case-by-case basis. A possible workaround could be to make the 'sensitiveness' of the processing dependent on the purpose of the processing (eg, the purpose of extracting biometric information). This way, the sensitivity can be determined in advance, but regulation then seems to depend on a catalogue of possible processing purposes. In addition, this may mean that the standards can only discipline cases of known problematic uses, but cannot protect against potentially harmful processing activities of a kind that is not yet known or not yet technically feasible.

In conclusion, the distinction between data and information seems to me to be a promising way to address the issue in a more structural way than when the issue of re-use of personal data is approached in terms of individual rights. However, this structural approach should be modified to allow for both a more systematic and a more flexible way of identifying potential improper re-uses in order to keep pace with technological innovation.

³²Elisa Orrù, 'Minimum Harm by Design. Reworking Privacy by Design to Mitigate the Risks of Surveillance,' in *Computers, Privacy and Data Protection: Invisibilities & Infrastructures*, ed. Ronald Leenes u. a. (Dordrecht: Springer, 2017), 129.

IV. A Contextual Approach: Information Flows Must Respect Contextual Norms

A third line of argument focuses on the contexts in which information flows, and sees potential violations of users' rights whenever information, as a result of secondary processing, appears in contexts different from the one in which it was originally shared.

Brown et al.,³³ for instance, discuss privacy issues raised by the use of training data for language models. They argue that current techniques for implementing privacy in language models are inadequate because they rely on misleading assumptions about the nature of information and information flows. Such techniques cannot avoid eg, the identification of members of the training data through adversarial attacks.³⁴

The most common approaches for privacy preserving language models either aim to remove all private information from the data used (data sanitisation) or use algorithms that do not memorise private information (differential privacy). These methods fail to provide effective privacy because they both assume that private information can be formally defined, easily identified, and isolated from other information. On the contrary, the authors argue that whether a piece of information is 'private' or not is not a property of the data itself, but depends on the context of its use. This means, among other things, that 'an appropriately named "privacy-preserving" LM [language model] should guarantee that a user's data cannot ever appear (or be inferable) outside the context they originally expected it to appear [...] – an ability that cannot be achieved without a deep understanding of the context in which the private information is produced, used, and shared.'³⁵

For their argument, Brown et al. explicitly draw on the concept of privacy as contextual integrity, as formulated by the US mathematician and philosopher Helen Nissenbaum.³⁶

Nissenbaum's conceptualisation is based on the sociologically inspired idea that individuals do not simply interact with each other as individuals, but always within specific social spheres or contexts that pre-shape their roles and expectations. This is also true of privacy expectations: these are not intersubjectively determined each time we exchange information, but rather are predetermined according to the context in which we interact and the roles we embody (ie, doctor/patient, teacher/student, friend/friend, employer/employee, vendor/customer, parent/child, etc). Each context is characterised by specific privacy

³³ Brown et al., 'What Does it Mean for a Language Model to Preserve Privacy?'

³⁴ *ibid* 3.

³⁵ *ibid* 2.

³⁶ Helen Nissenbaum, 'Privacy as Contextual Integrity Symposium – Technology, Values, and the Justice System,' *Washington Law Review*, no. 79 (2004): 119–58; Nissenbaum, *Privacy in Context*.

norms that define what content and topics are appropriate and what constraints should apply to the transfer of information (ie, confidentiality, reciprocity, asymmetry, etc).³⁷

According to Nissenbaum 'many novel information flows are disruptive not because they contravene explicit norms, but because they open up previously impossible (possibly unimaginable) flows. In these instances, consternation follows because flows are unprecedented, may or may not expose new vulnerabilities and hazards.'³⁸ In relation to the AI applications that Brown et al. are concerned with, and against the background of Nissenbaum's theory, it can be argued that LMs create new information flows that disrupt previously existing contextual boundaries, and thus bring about the transfer of information that was intended for a specific context governed by specific informational norms to new contexts that are governed by different norms or in which no established informational norms yet exist.

On the one hand, a strength of an approach inspired by the contextual integrity framework to the cases of reuses of publicly available data seems to lie in the provision of a general rule for identifying privacy violations: this would then consist in transposing information into a new context.³⁹ It thus provides a way of overcoming the case-by-case approach of the above accounts. Nissenbaum's contextual approach adapted to secondary processing of publicly available data thus provides a normative basis for identifying prima facie cases of privacy violations that require additional caution in (or prohibition of) further processing. The precise identification of the source and destination contexts and their contextual norms would then be a matter of case-by-case assessment, but at least a general norm, not linked to a catalogue of specific uses, can be formulated.

On the other hand, Nissenbaum's theory has been criticised for relying too much on the analysis of existing norms.⁴⁰ Deriving from this general criticism, two main reproaches have been formulated against her theory: first, that it is normatively weak, and second, that it is (admittedly) conservative. The first line of criticism, in line with the prohibition – philosophically canonical since David Hume, but not uncontroversial – of deriving prescriptions from descriptions,⁴¹ sees a problem in the fact that Nissenbaum's approach derives contextual norms

³⁷ Nissenbaum, 'Respecting Context to Protect Privacy,' 838–839; Nissenbaum, *Privacy in Context*, 140–47.

³⁸ Nissenbaum, 'Respecting Context to Protect Privacy,' 841.

³⁹ This can be seen as an adaptation of Nissenbaum's account of privacy violations that occur when contextual norms are violated. See Nissenbaum, *Privacy in Context*.

⁴⁰ See as paradigmatic critiques the ones of James B. Rule and, more sympathetic to Nissenbaum's account, Marcel Becker: James B Rule, 'Contextual Integrity and Its Discontents: A Critique of Helen Nissenbaum's Normative Arguments,' *Policy & Internet* 11, no. 3 (2019): 260–79, doi.org/10.1002/poi3.215; Marcel Becker, 'Privacy in the Digital Age: Comparing and Contrasting Individual versus Social Approaches towards Privacy,' *Ethics and Information Technology* 21, no. 4 (December 2019): 307–17, doi.org/10.1007/s10676-019-09508-z.

⁴¹ David Hume, *A Treatise of Human Nature* (Oxford: Clarendon Press, 2011).

from existing social practices. The second argument is that her account is too resilient to change, including the legitimate replacement of old contextual norms with new ones. Moreover, Nissenbaum's theory would assume social homogeneity and universal agreement about what norms count in a given context, whereas in fact these are highly contested.

Overall, a contextual approach presents the advantage of providing a systematic approach to dealing with the re-processing of publicly available data. In my view, however, it should allow for a more open and active attitude towards legitimate modifications of contextual norms. Although technological innovation does have a disruptive effect on existing contextual norms, it is not the only factor driving changes in privacy norms. Changes in contextual norms, as I will argue below, can also be triggered by a need for greater protection and can be enacted through participatory, democratic processes that can ultimately lead to new regulations.

V. Which Approach Best Protects Privacy when it Comes to Publicly Available Data?

The examples of secondary processing of publicly available data discussed above show once again that in the era of big data and artificial intelligence, concepts of privacy and data protection that focus solely on the (atomistic) individual are less suited than ever to ensuring effective privacy and data protection. How can individuals alone deal with the huge asymmetries in power and technical means that separate them from public and private organisations such as government agencies and big tech companies? How can they track how data about them are used and reused once they are 'out there', and understand the consequences of these practices in order to legally enforce their rights? As Solove notes:

Rights can give people a small amount of power in a few isolated instances, but this power is too fragmented and haphazard to have a meaningful impact on protecting privacy. Ultimately, rights are at most capable of being a supporting actor, a small component in a much larger architecture.⁴²

If individual rights are only a small part of a much-needed more systematic approach, what principles should shape the 'larger architecture' to make privacy protection effective?

Distinguishing between data and information and shifting the focus from the origin of the 'data' to the uses made of them and the information extracted from them, as suggested by Albers, seems to be a reasonable way to overcome an overly individualistic approach. Shifting the focus from the origin to the use of data allows to enrich the individual and 'negative' rights perspective with a more

⁴² Solove, 'The Limitations of Privacy Rights,' 978.

systematic and 'positive' component. Crucial questions such as the 'sensitiveness' of the data/information no longer focus only on the 'source' of the data, but also on the flows and uses of the information. In other words, the perspective becomes less individualistic and more structural. Indeed, ensuring that data are handled appropriately is a task that can only be fulfilled if public authorities create the structural conditions that allow data to flow appropriately (and individual rights to be enforced). And yet, as we have seen, this perspective alone relies too much on case-by-case considerations to determine when a privacy violation is likely to occur or has occurred, and is not flexible enough to deal with new kinds of re-processing.⁴³

A potential mitigation of this shortcoming lies in combining the structural perspective with contextual approaches: according to this combined approach, privacy violations are likely to occur whenever data are processed in a context different from the one in which they were originally shared. The change or disruption of the context could then serve as a first criterion to identify situations where additional precautions and restrictions need to be applied to data processing. Under this approach, data processors who wish to re-use publicly available data would then be responsible for convincingly demonstrating that the further processing is in line with the informational norms that governed the original context. At this more concrete level, considerations based on the distinction between data and information play an important role in determining whether the uses in the new contexts are consistent with the informational norms of the original context.

A contextual approach also seems to be particularly appropriate to face the challenges of the digital and AI age. Indeed, I propose to consider the main characteristic of the digital word, of big data and even more of AI applications, as their disruption of contexts. What makes big data and AI applications so distinctive, powerful and, for some, threatening, is that they make it possible to connect different social, epistemic and interactional contexts with an exponential increase in speed and ease.⁴⁴ It may be that the first legal conceptualisations of a right to privacy were prompted by a similar challenge to the disruption of contexts. When Samuel Warren and Louis Brandeis published their famous article 'The Right of Privacy' in 1890, it was probably the diffusion of photography and the subsequent possibility of disseminating images taken in a private context

⁴³ See section III above.

⁴⁴ In a similar vein, Nissenbaum notes: 'The challenge of privacy online is not that the venue is distinct and different, or that privacy requirements are distinct and different, but that mediation by the Net leads to disruptions in the capture, analysis, and dissemination of information as we act, interact, and transact online.' H Nissenbaum, 'A Contextual Approach to Privacy Online,' *Daedalus, the Journal of the American Academy of Arts & Sciences* 140, no. 4 (2011): 38. Also the notion of 'contextual gaps' points into a similar direction, see Gordon Hull, Heather Richter Lipford, and Celine Latulipe, 'Contextual Gaps: Privacy Issues on Facebook,' *Ethics and Information Technology* 13, no. 4 (December 2011): 289–302, doi.org/10.1007/s10676-010-9224-8.

to a wide public that triggered their reflections.⁴⁵ Similarly, data produced and information shared in specific social contexts can now easily be transposed, used and disseminated in contexts of which the people involved were not even aware. This disrupts not only the contexts of information sharing and their norms, but also the very reliability of the separation of contexts that has been at the core of social interaction.⁴⁶

Finally, a combination of the structural and contextual approaches makes it possible to address the critique of conservatism that has been advanced to purely contextual approaches.⁴⁷ By emphasising the need for regulatory powers to intervene and shape information flows, a structural approach points to one of the most important actors able to actively influence and change informational norms, namely legislative powers. This may sound like an overly legalistic and law-centred perspective. Adopting this perspective, however, does not mean neglecting the importance of other sources of informational norms or the fact that legal norms are not the only kind of societal norms, which are also shaped by morality, conventions, religion, culture, and so on. Moreover, the legal sphere is not rigidly separated from the other normative spheres, and the legislative process is, at its best, sensitive to challenges and demands for change from society. The GDPR, despite all its limitations, can be seen as a positive example of legal norm-making that has been able to partially reshape contextual norms in a way that effectively responds to technological changes and the discomfort caused by emerging practices that disrupt existing information contexts and norms. In addition, new laws influence individuals' perceptions and attitudes and can thus trigger or facilitate a shift in non-legal contextual norms. For example, even if notice and consent policies are mostly designed to nudge users into accepting cookies, it is still important for shaping moral, cultural and societal norms that users know that the use of cookies should not legally be the default option and should only be allowed with explicit consent.

VI. Conclusion

This chapter addressed the question as to whether publicly available (personal) data can be used for arbitrary purposes simply by virtue of their public accessibility.

⁴⁵ Samuel D Warren and Louis D Brandeis, 'The Right to Privacy [the Implicit Made Explicit]'; in *Philosophical Dimensions of Privacy: An Anthology*, eds. Ferdinand David Schoeman (Cambridge: Cambridge University Press, 1984), 75–103; Amy Gajda, *What if Samuel D. Warren Hadn't Married a Senator's Daughter?: Uncovering the Press Coverage That Led to the Right to Privacy*, vol. 6, Illinois Public Law and Legal Theory Research Papers Series 7, 2007.

⁴⁶ Erving Goffman, *The Presentation of Self in Everyday Life* (New York: Anchor, 1959); Karl Lenz, 'The Presentation of Self in Everyday Life', in *Goffman-Handbuch: Leben – Werk – Wirkung*, eds. Karl Lenz and Robert Hettlage (Stuttgart: J.B. Metzler, 2022), 267–74, doi.org/10.1007/978-3-476-05871-3_37; Nissenbaum, *Privacy in Context*.

⁴⁷ For the criticism see section IV above.

It analysed three types of arguments used in the literature to answer this question: individual rights arguments centred on the notion of consent, structural considerations based on the distinction between data and information, and finally contextual conceptions focusing on the norms that regulate information flows in different societal contexts. The chapter argued that a combination of the structural and contextual approaches is best suited to address the challenges posed by digital technologies, big data and AI systems. Indeed, it argued that the core characteristic of these techniques is their ability not only to disrupt the informational context, but also to render obsolete the reliability of distinguishing between different informational contexts. Agreeing with Solove's assertion that 'effective privacy protection involves not just facilitating individual control, but also bringing the collection, processing, and transfer of personal data *under control*',⁴⁸ I have argued for a modification of Nissenbaum's contextual approach, which incorporates elements of Albers's structural approach and her distinction between data and information and emphasises the positive role of law in shaping and directing contextual norms.

The systematising effort done in this article is meant to provide the basis for further work aiming at translating the conceptual background into practical ethical guidance for re-processing of publicly available data, especially but not exclusively for research purposes.

References

- Albers, Marion. *Informationelle Selbstbestimmung*. Baden-Baden: Nomos, 2005. doi.org/10.5771/9783845258638.
- . 'Realizing the Complexity of Data Protection.' In *Reloading Data Protection*, edited by Serge Gutwirth, Ronald Leenes, and Paul de Hert, 213–35. Dordrecht: Springer, 2014. doi.org/10.1007/978-94-007-7540-4_11.
- . 'Umgang mit personenbezogenen Informationen und Daten.' In *Grundlagen des Verwaltungsrechts*, edited by Andreas Voßkuhle, Martin Eifert, Christoph Möllers, Wolfgang Hoffmann-Riem, and Eberhard Schmidt-Aßmann, 2: § 22. München: C.H. Beck, 2006.
- Brown, Hannah, Katherine Lee, Fatemehsadat Mireshghallah, Reza Shokri, and Florian Tramèr. 'What Does it Mean for a Language Model to Preserve Privacy?'. In *Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*, 2280–92. FAccT '22. New York, NY, US: Association for Computing Machinery, 2022. doi.org/10.1145/3531146.3534642.
- Coppersmith, Glen, Kim Ngo, Ryan Leary, and Anthony Wood. 'Exploratory Analysis of Social Media Prior to a Suicide Attempt'. In *Proceedings of the Third Workshop on Computational Linguistics and Clinical Psychology*, 106–17. San Diego, CA, US: Association for Computational Linguistics, 2016. doi.org/10.18653/v1/W16-0311.

⁴⁸ Solove, 'The Limitations of Privacy Rights,' 975, emphasis added.

- criticalai. 'CRITICAL AI: Adapting College Writing for the Age of Large Language Models Such as ChatGPT: Some Next Steps for Educators.' 17 January 2023. criticalai.org/2023/01/17/critical-ai-adapting-college-writing-for-the-age-of-large-language-models-such-as-chatgpt-some-next-steps-for-educators/.
- De Choudhury, Munmun, Scott Counts, and Eric Horvitz. 'Predicting Postpartum Changes in Emotion and Behavior via Social Media.' In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 3267–76. CHI '13. New York, NY, US: Association for Computing Machinery, 2013. doi.org/10.1145/2470654.2466447.
- Gajda, Amy. *What if Samuel D. Warren Hadn't Married a Senator's Daughter?: Uncovering the Press Coverage That Led to the Right to Privacy*. Vol. 6. Illinois Public Law and Legal Theory Research Papers Series 7, 2007.
- Goffman, Erving. *The Presentation of Self in Everyday Life*. New York: Anchor, 1959.
- Hornung, Gerrit, and Carolin Gilga. 'Einmal öffentlich – Für immer schutzlos?' *Computer und Recht* 36, no. 6 (June 2020): 367–79. doi.org/10.9785/cr-2020-360609.
- 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. doi.org/10.1007/s10676-010-9224-8.
- Hume, David, *A Treatise of Human Nature*. Oxford: Clarendon Press, 2011.
- Jobin, Anna, Marcello Ienca, and Effy Vayena. 'The Global Landscape of AI Ethics Guidelines.' *Nature Machine Intelligence* 1, no. 9 (September 2019): 389–99. doi.org/10.1038/s42256-019-0088-2.
- Jol, Guusje, and Wyke Stommel. 'Ethical Considerations of Secondary Data Use: What about Informed Consent?' *Dutch Journal of Applied Linguistics* 5, no. 2 (January 2016): 180–95. doi.org/10.1075/dujal.5.2.06jol.
- Lenz, Karl. 'The Presentation of Self in Everyday Life.' In *Goffman-Handbuch: Leben – Werk – Wirkung*, edited by Karl Lenz and Robert Hettlage, 267–74. Stuttgart: J.B. Metzler, 2022. doi.org/10.1007/978-3-476-05871-3_37.
- Nicholas, Jennifer, Sandersan Onie, and Mark E Larsen. 'Ethics and Privacy in Social Media Research for Mental Health.' *Current Psychiatry Reports* 22, no. 12 (November 2020): 84. doi.org/10.1007/s11920-020-01205-9.
- Nissenbaum, Helen. 'A Contextual Approach to Privacy Online.' *Daedalus, the Journal of the American Academy of Arts & Sciences* 140, no. 4 (2011): 32–48.
- . 'Privacy as Contextual Integrity Symposium – Technology, Values, and the Justice System.' *Washington Law Review* 79 (2004): 119–58.
- . 'Protecting Privacy in an Information Age: The Problem of Privacy in Public.' *Law and Philosophy* 17, nos. 5/6 (1998): 559–96.
- . *Privacy in Context: Technology, Policy, and the Integrity of Social Life*. Stanford: Stanford Law Books, 2010.
- . 'Respecting Context to Protect Privacy: Why Meaning Matters.' *Science and Engineering Ethics*, 24 (2018): 831–852. doi.org/10.1007/s11948-015-9674-9.
- Orrù, Elisa. 'Minimum Harm by Design. Reworking Privacy by Design to Mitigate the Risks of Surveillance.' In *Computers, Privacy and Data Protection: Invisibilities & Infrastructures*, edited by Ronald Leenes, Rosamunde Van Brakel, Serge Gutwirth, and Paul De Hert, 107–37. Dordrecht: Springer, 2017.
- Ploug, Thomas. 'The Right Not to Be Subjected to AI Profiling Based on Publicly Available Data—Privacy and the Exceptionalism of AI Profiling.' *Philosophy and Technology* 36, no. 1 (2023): 1–22. doi.org/10.1007/s13347-023-00616-9.

- Ravn, Signe, Ashley Barnwell, and Barbara Barbosa Neves. 'What Is 'Publicly Available Data'? Exploring Blurred Public-Private Boundaries and Ethical Practices Through a Case Study on Instagram.' *Journal of Empirical Research on Human Research Ethics* 15, no. 1-2 (February 2020): 40-45. doi.org/10.1177/1556264619850736.
- Rule, James B. 'Contextual Integrity and Its Discontents: A Critique of Helen Nissenbaum's Normative Arguments.' *Policy & Internet* 11, no. 3 (2019): 260-79. doi.org/10.1002/poi3.215.
- Schindler, Stephan, and Gerrit Hornung. 'Datenschutz bei der biometrischen Gesichtserkennung.' *Datenschutz und Datensicherheit - DuD* 45, no. 8 (August 2021): 515-21. doi.org/10.1007/s11623-021-1482-6.
- Solove, Daniel J. 'The Limitations of Privacy Rights.' *Notre Dame Law Review* 98 (2023): 975-1036. dx.doi.org/10.2139/ssrn.4024790.
- Stommel, Wyke, and Lynn de Rijk. 'Ethical Approval: None Sought. How Discourse Analysts Report Ethical Issues around Publicly Available Online Data.' *Research Ethics* 17, no. 3 (July 2021): 275-97. doi.org/10.1177/1747016120988767.
- Susser, Daniel. 'Notice After Notice-and-Consent: Why Privacy Disclosures Are Valuable Even if Consent Frameworks Aren't.' *Journal of Information Policy* 9 (2019): 148-73. doi.org/10.5325/jinfopoli.9.2019.0148.
- Taylor, Joanna, and Claudia Pagliari. 'Mining Social Media Data: How Are Research Sponsors and Researchers Addressing the Ethical Challenges?.' *Research Ethics* 14, no. 2 (October 2017). doi.org/10.1177/1747016117738559.
- UNESCO and EQUALS Skills Coalition. 'I'd Blush if I Could: Closing Gender Divides in Digital Skills through Education', 2019. unesdoc.unesco.org/ark:/48223/pf0000367416.page=1.
- Warren, Samuel D, and Louis D Brandeis. 'The Right to Privacy [the Implicit Made Explicit]'. In *Philosophical Dimensions of Privacy: An Anthology*, edited by Ferdinand David Schoeman, 75-103. Cambridge: Cambridge University Press, 1984.

