Telecommunications Policy 47 (2023) 102446

Contents lists available at ScienceDirect

Telecommunications Policy

journal homepage: www.elsevier.com/locate/telpol

An unrealistic and undesirable alternative to the right to be forgotten

Björn Lundgren [a,b,c,*]

[a] Department of Philosophy and Religious Studies, Utrecht University, Utrecht, the Netherlands
[b] Institute for Futures Studies, Stockholm, Sweden
[c] Department of Philosophy, Stockholm University, Stockholm, Sweden

ARTICLE INFO

Keywords:
Privacy
The right to privacy
The right to be forgotten
Opacity
Transparency
Manipulation
Social norms

ABSTRACT

In their "*Così fan tutte*: A better approach than the right to be forgotten", Martha Garcia-Murillo and Ian MacInnes (2018) present an alternative to the right to be forgotten. It mainly consists of changing social norms to a more open and transparent society and improving anti-discrimination laws. In this reply, I challenge their suggestions both empirically and normatively, arguing that they cannot achieve what they set-out; that their analysis of privacy and the right to be forgotten misses important aspects that go beyond discrimination; and that we have a shared interest in allowing individuals to live a partly opaque life.

1. Introduction

In their "*Così fan tutte*: A better approach than the right to be forgotten", Martha Garcia-Murillo and Ian MacInnes present a number of arguments against the right to be forgotten while also promoting an alternative. In this brief reply, I will focus on challenging their proposed alternative, which depends on changing societal norms, and improving upon anti-discrimination laws.[1] Simply put, there is no need to forget if people do not discriminate against us. Their basic idea is that in a more accepting society the right to be forgotten is unnecessary. However, their social agenda does not stop with anti-discriminating, they also want to change our information distribution norms:

> We believe that a more open, transparent, and forgiving society will allow people to communicate more freely without having to be concerned about the implications or hide behind anonymity to express an unpopular opinion. Garcia-Murillo and MacInnes (2018, p. 236)

I will argue that their solution is unrealistic and undesirable. I will begin by briefly commenting on the empirical limitations of their proposal, first, in regard to what anti-discrimination legalization can achieve. Second, I will turn to discuss the problem of changing social norms, and why biases are unavoidable. After these empirical critiques I will turn to two normative critiques. First, I will argue

* Department of Philosophy and Religious Studies, Utrecht University, Utrecht, the Netherlands.
E-mail addresses: b.a.lundgren@uu.nl, bjorn.lundgren@iffs.se, bjorn.lundgren@philosophy.su.se.

[1] They also recognize the importance of protecting individuals against manipulation. However, I fail to see what their proposal regarding this worry is except for implementing "standards of information quality" (Garcia-Murillo and MacInnes, 2018, p. 234), which does not do nearly enough to recognize the problem of data-based manipulation (see, e.g., Lundgren, 2020a; Véliz, 2020). Given that their proposal is not substantial enough to address the challenges that we are facing due to online-manipulation, I will set this issue aside and focus on the more developed part of their solution.

https://doi.org/10.1016/j.telpol.2022.102446
Received 10 June 2022; Received in revised form 14 September 2022; Accepted 14 September 2022
Available online 8 October 2022

that the right to be forgotten cannot solely be about avoiding discrimination. Indeed, that is a very flawed view of privacy. Second, I will argue that part of their ideal society—openness and transparency—is not desirable. While we should desire openness and transparency on institutional levels, there needs to be a modicum of privacy available on an individual level. Finally, I will end the article by summing up my arguments and briefly commenting on how my criticism should be weighed against Garcia-Murillo and MacInnes' criticism of the right to be forgotten.

2. Is Garcia-Murillo and MacInnes' proposal realistic?

As noted in the introduction, Garcia-Murillo and MacInnes think that we should discard the right to be forgotten and replace it with other means that better address the problems that the right to be forgotten aims to resolve (according to their analysis). They think that we can fix some of the problems by improving upon anti-discrimination legislation. Indeed, one of the problems that the right to be forgotten aims to eliminate is the risk that individuals are adversely judged by their peers because of events from their past. While I agree with Garcia-Murillo and MacInnes' call for improving legal protections against discrimination and aiming to try resolve some of the policy aims of a right to be forgotten by other means, I disagree with them that their proposal is a realistic *replacement*.

It is perhaps ironic that they propose legal remedies, given that their critique of the right to be forgotten is partly based a complaint that the current legal practices—as they apply to the right to be forgotten—are "inconsistent and unreliable" Garcia-Murillo and MacInnes (2018, p. 232). Although they may be correct in their critique, it is not clear that legal solutions to discrimination are better off. Indeed, *pace* Garcia-Murillo and MacInnes, legal solutions to discrimination are notoriously difficult. Consider, for example, recent evidence "that prestige is associated with a decreased likelihood of being found liable" in cases of employment discrimination (McDonnell & King, 2018, p. 61), which illustrates the severe challenges of a legislative solution to anti-discrimination. Garcia-Murillo and MacInnes attempt to pre-empt this kind of argument by saying that claims of naiveite and unrealistic expectations are disproven by "many previously successful changes in social behavior and is, thus, simply incorrect" Garcia-Murillo and MacInnes (2018, p. 236).

However, even if there is room for improvement in legal protection in the area of anti-discrimination, that does not mean that progress has no upper bound; in fact, presuming otherwise is just a fallacy. Simply put, there are theoretical limitations to solving discrimination by legal measures. The problem has to do with the lack of evidence in singular cases. For example, while a singular applicant for a job cannot—upon receipt of a standard rejection letter—know whether they were rejected on the basis of discriminatory practices, a weaker CV, or bad fortune, we know that "extensive ethnic labor market discrimination [persists] across countries" (Bursell et al., 2021). That is, whether a person is discriminated against is—to some extent—a matter of *intentions*. While we can, through experiments, statically determine that there is extensive discrimination in the labor market *overall*, there are limitations to providing evidence in *singular* cases. Because discrimination is intentional and intentions do not, as per their definition, provide us with proof, it is challenging to prove discrimination in singular cases. Simply put, "crimes of the mind" do not necessarily provide evidence that can be found in the world. Clearly, the evidentiary limitations in singular cases cannot be overcome and solved by legislation because that would require an ability to measure individuals' past intentions and such technologies are not readily available for policy purposes. Moreover, even if we could measure people's intentions, the evidence we need would be *past* intentions, not *current* intentions (although the latter could provide some indirect evidence for the former). Thus, although it should be recognized that there is potential for progress on an institutional level, we also need to recognize that as long as singular events cannot be jointly related to one perpetrator (e.g., a singular company), there are clear limitations as to what anti-discrimination laws can do to resolve the issue of discrimination.

Of course, Garcia-Murillo and MacInnes do not think that legislation alone can solve the problems available. They also want to complement the legislation with norm changes. They believe that changing our social norms would allow us to live our lives in a society better adapted to more transparent information practices. However, changing social norms is arguably not easier than legislation. If we, following Bicchieri (2017), understand social norms as second-order beliefs—more specifically, beliefs about what others think is the correct thing to do in a given situation—then there are many cases in which social norms dominate even when most people think that the norms are wrong; indeed, Bicchieri (2017) provides many such examples. Of course, Garcia-Murillo and MacInnes are arguing that norms are already changing, that is, that young people are more comfortable sharing information online. Whether that is relevant depends partly on a normative argument, which I will address in the next section; for now, all that needs to be said is that the evidence that Garcia-Murillo and MacInnes consider leaves something to be desired. There are two problems. First, the evidence they supply to support the thesis that the "[y]outh today are comfortable expressing themselves online" Garcia-Murillo and MacInnes (2018, p. 233) is rested on a singular anecdote about a young women's blog practices. Second, even if there is a norm change—and I believe we should grant that there is—that does not provide evidence for future changes, or that the norms under discussion are universal and uniform across the youth generation. This raises questions about whether the proposal is realistic (i.e., norms haven't changed as universally and uniformly as Garcia-Murillo and MacInnes seem to think) as well as normative questions (i.e., if we adopt policy fitted to those with a preference for sharing private information, then we are ignoring those with contrary preferences).

Lastly, there is an overarching reason to think that humans will never reach social perfection and that we may therefore need regulations that reflect our limitations. Simply put, humans are unavoidably biased. While we may diminish some biases, our cognitive system is pattern-seeking, which is highly beneficial but also means that we unavoidably find patterns where there are none, or simply identify the wrong patterns. While biases can be diminished, we cannot avoid our pattern-seeking nature completely and the risk of bias that follows from it.

3. Does Garcia-Murillo and MacInnes' proposal address the normative challenges and is it desirable?

I will now turn to the more normative aspects of their arguments. The goal of this section is to argue that the value of privacy and the right to be forgotten goes beyond discrimination (and even beyond manipulation, cf. fn. 1) and that the changes that Garcia-Murillo and MacInnes envision are not desirable. Indeed, while some hold that privacy is about freedom of judgement from others (Johnson, 1989)—an analysis that Garcia-Murillo and MacInnes arguably could have attempted to rely on—we can easily see the limitation of such an idea by considering the following example:

> even married couples whose sex-lives are normal (whatever that is), and so who have nothing to be ashamed of, by even the most conventional standards, and certainly nothing to be blackmailed about, do not want their bedrooms bugged. (Rachels, 1975, p. 325)

While Rachels' example is meant as a counterexample to the idea of analyzing the concept of (the right to) privacy in terms of its instrumental function, it also serves our present purpose. Indeed, even if we set aside the risks of discrimination and manipulation, we may still desire that certain aspects of our lives remain private or—in the case of their wide distribution—that they could be forgotten.

Simply put, the idea that discrimination (and manipulation, cf. fn. 1) are all that there is about privacy seems to run contrary to a long tradition of understanding privacy as important—or even necessary—for other aspects of our lives, such as friendship or intimacy (see, e.g., Fried, 1970; Gerstein, 1978; Rachels, 1975). That is, beyond the risk of discrimination and manipulation, privacy can serve an important function in enabling and maintaining intimate relationships (because if they are no longer private, then it may be difficult to maintain that intimate part of the relationship) and in regulating our social relations—differentiating between lovers, friends, acquaintances, colleagues, partners, strangers, fiduciary relationships, and other forms of social relations. The idea that privacy is important for our social relationships has often been rested on an idea of the right to privacy as conceptualized in terms of control. That idea then naturally lends itself to the idea that control over information matters for achieving different forms of social relationship.[2] This idea has recently been defended by Marmor (2015, pp. 3–4), who argues that the "right to privacy [is] 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." Following this line of thought the right to be forgotten can play an important role in ensuring that control. Of course, we need not rely on control views in general, nor on Marmor's view in particular (for critiques of Marmor, see, e.g. Lundgren, 2020a; 2020b; Munch, 2020). We can motivate the importance of a right to be forgotten as a means to ensure that we are not prejudged by other people, as a means of enabling us to form social relationships. Such a view need not be rested on a control view.[3]

Moreover, we can use the above example to ask another set of questions, which I will now turn to: Do we really want the norms that Garcia-Murillo and MacInnes promote? How open and transparent do we want our lives to be? It is illustrative to consider what it would mean to maximize Garcia-Murillo and MacInnes' desire for openness and transparency. That is, although Garcia-Murillo and MacInnes do not argue for *absolute* openness and transparency, I believe that I can show the limits of their proposal by considering the consequences of a society in which there is no informational opacity. Such a society is—of course—empirically unrealistic, but is it desirable?

An illustrative example that is often used in cases such as this is the example of living in a panopticon, in which we are all being monitored by everyone else. (*N.B.*, in the original panopticon—set in a prison—the prisoners were not only being watched by each other, but also by guards, whom they could not see. In a completely transparent world, there would of course be nothing similar to the hidden guards. However, complete transparency does not mean that everything is stored. Hence, our ability to know whether we've been "surveilled" would only be possible if we also added the premise of absolute surveillance and data storage, which perhaps is in line with Garcia-Murillo and MacInnes' desire to make use of more data; see Garcia-Murillo and MacInnes, 2018, p. 232.)[4]

So what is the privacy-related problem with a panoptic world? Simply put, in such a world, there is no longer room to draw an information-based limit between friends and foe (or simply: non-friends); nor is there any room for one's private self. We do not need to consider graphical examples of bathroom visits to realize that a panoptic world is a world that shares features with Orwell's *1984*;

[2] For a modern critique of part of Rachels' (1975) original thesis, see Mooradian (2009). However, given that Mooradian's critical argument concerns the overall *collection* of personal information by services such as Facebook, that criticism is not very relevant for the given context, since what I am addressing above is information distributions between people and the problems that follow from that. In fact, Mooradian seems to think that Rachels' theory still is highly relevant in this context, since "[i]t offers insights into how information flows within social networking sites are related to the creation and maintenance of personal relationships" (Mooradian 2009, p. 173).

[3] Here it may be worthwhile to point out that Garcia-Murillo and MacInnes believe that there is an analytical connection between the right to be forgotten and the idea of the right to privacy as conceptualized in terms of control. Yet, this is simply false, as the argument above shows. Another way to motivate the right to forgotten is based on the framework of contextual integrity (see Nissenbaum, 2010 for a full account of her theory). A core idea of Nissenbaum's theory is that there are contextually sensitive information norms that hold across social spheres. Most importantly in this context, Nissenbaum does not think that the online context constitutes a distinction social sphere. Conversely, the information norms that should guide the right to privacy online are based on offline norms (Nissenbaum, 2011). Based on this it is easy to see how her theory could provide a motivation for a right to be forgotten. In an offline context, people forget and documents are not as easily accessible. Consider the following example: search engines make old news articles easily accessible, but finding old news articles in the offline world requires more of an effort. Thus, the right to be forgotten can be seen as a right that mediates the online world to the offline realities in accordance with Nissenbaum's theory.

[4] The example of the panopticon has a rich history in various debates, including punishment and privacy. For a detailed description of the panopticon, see Foucault (1995[1977]), who also discusses it in relation to surveillance and power relations. For a recent privacy-related discussion, see Marmor (2015).

instead of censorship and external thought control, we would have self-censorship and internal thought control (arguably, the thought control could also be external, granted that the norm changes that Garcia-Murillo and MacInnes envision would not fully mature). This keenly illustrates that despite the value of a more open society, transparency on an individual level is also problematic. Simply put, if we believe that there is sometimes value in solitude (individually or in smaller groups), then in order to live an autonomous life we must sometimes be able to restrict others' access to ourselves. If so, we should conclude that opacity is important to live an autonomous life.

It is possible to respond to the above argument by questioning the need for a right to be forgotten. Indeed, someone might think that we can live a fully autonomous life without the right to be forgotten as long as there are sufficient data and privacy right protections. While I agree that this is *possible*, it is only possible as long as we treat data and privacy right protections as absolute. As soon as we engage in certain forms of trade-offs, the right to be forgotten becomes a relevant means to mediate these trade-offs. For example, there may be an important news article that all-things-considered ought to be published, despite infringing upon someone's right to privacy.[5] However, cases in which the right to privacy is—for example—overridden by other concerns, are often contextually sensitive. In the given example, the news value of a certain event may be extremely time-sensitive, which illustrates quite nicely the role of the right to be forgotten. Indeed, while an infringement of someone's right to privacy may, all-things-considered, be motivated because of, say, a public interest in a given context and a certain time, the ethical considerations may have changed substantially 10 years later. It is when such considerations are altered that the right to be forgotten serves as an important means to safeguard the balance between data and privacy right protections and other considerations (such as public interest). That is, 10 years later, the public interest may no longer outweigh privacy considerations in such a way that motivates the easy online-access provided by the current online praxis.[6] Moreover, even if we treat data and privacy right protections as absolute, there is also a risk of mistakes. The right to be forgotten can serve as a means to remedy such mistakes. Thus, we can do without the right to be forgotten but only if data protection and privacy rights are absolute and if no mistakes in managing these rights are ever made.

Alternatively, some may think that it is possible to live a fully autonomous life while being fully surveilled. However, as far as I know there are no examples that support such a thesis (despite what we can learn about surveillance and autonomy within, for example, prisons). Moreover, to further illustrate my point let us reconsider Rachels' example as follows: While there are groups of people who happily share private encounters, which we normally take to be intimate, with a broad audience (through TV, or social media), we can note that these are exceptions to the normal practice and, more importantly, we can question whether these encounters are truly intimate.

Another potential criticism of my argument would be to question the extent to which we can draw conclusions from an example that, admittedly, is exaggerated. However, if we follow the arguably vague proposal of Garcia-Murillo and MacInnes, then—given that my arguments so far are correct—*some* situations that require opacity to protect our privacy will be transparent.[7] While we can argue that there may be a trade-off between the benefits of their proposal and the worries I raise here, I would note that the current costs of the data-based economy would override any such benefits (see, e.g., Lundgren, 2020a; Véliz, 2020).[8]

Lastly, we should not presume that more transparency means that we would have a more equal or egalitarian distribution of informational resources. Just as the online world is functioning right now, people, companies, organizations, or states with means beyond ordinary individuals will always be in a better position to make use of big data troves, simply because they require computational power to make use of the data, which in turn requires competence and substantial resources. However, because of the unequal distribution of computational resources, increased transparency would arguably increase the risk of manipulation, power asymmetries, and informational injustice. This would, in particular, benefit larger organizations at the cost of individuals. It also important to note that these worries go beyond purely privacy concerns and challenge our basic democratic foundations. Given that information can be used to manipulate and influence people on a large scale and that elections are often determined by small margins, these

[5] Privacy trade-offs are also a central part of machine (or AI) ethics (see Lundgren, 2021a).

[6] Thanks to an anonymous reviewer for challenging me on this.

[7] An anonymous reviewer raised the issue of what standard of normative judgement that is supposed to guide us when we address the question of under which conditions opacity is appropriate. As I note above, there will be trade-offs between privacy and other considerations, but what is at stake here is whether I can show that privacy will be threatened by transparency. To show that let us consider what I call "our ability to be anonymous" (Lundgren, 2020a). Anonymity, or the ability to maintain it, can be seen as a safeguard for privacy. Consider, for example, how different privacy-protection technologies such as differential privacy (see, e.g., Dwork & Roth, 2014) protect privacy by ensuring individuals' anonymity. The problem, as I argue in Lundgren (2020a), is that the ability to be anonymous is reduced whenever data or information about an individual is linked (even if it is not linked to the individual, which it is about, and even if that individual's anonymity is retained). This means that more data and information transparency will reduce our ability to be anonymous (even if that data/information is not privacy-sensitive as such), which in turn will also minimize a safeguard for privacy. Thus, privacy will be at risk in a more information-transparent society. In Lundgren (2020c; 2021b; 2021c), I suggest that the right to privacy can be infringed or violated by substantial risks. To see that this is the case, consider, for example, a cloud-service that sells secure storage of private data. Suppose that this cloud service has made your private data available in an unencrypted form, but that it is never actually accessed. Granted that no-one accesses your private data, we should say that your privacy is preserved, but at the same time the cloud-service has clearly violated your right to privacy.

[8] Of course, there may be benefits with increased transparency, even for privacy. Indeed, as noted by Brin (1998) transparency can also allow us to hold privacy-violators accountable. However, what we are facing today is perhaps best described as one-directional transparency. That is, individuals are transparent towards big data companies, but not the other way around.

manipulations need not be very effective on an individual level; as long as they can affect a small minority, they can sway an election (see, e.g., van den Hoven, 1997; Lundgren, 2020a; Véliz, 2020).[9]

4. Summation and concluding comments

In conclusion, I have argued that Garcia-Murillo and MacInnes' proposal is empirically unrealistic (there is clear limit to what legislation can do and how much we can change our social norms or our cognitive biases). Moreover, Garcia-Murillo and MacInnes' normative ideal misses aspects of privacy that remain important (solitude for ourselves or those we selectively choose be private with). The idea of transparency cuts against those basic values of privacy and our ability to lead autonomous lives (without self-censorship), which in turn would also increase the risk of manipulation.

Lastly, if we grant that some of Garcia-Murillo and MacInnes' critique holds, then one may wonder whether their ore my proposal is worse all things considered. Given my aim to keep the article brief, I will only note that part of Garcia-Murillo and MacInnes's critiques are about the practical application of the right to be forgotten, and such practices could be changed without having to scrap the principle in its entirety. Moreover, in cases when the right to be forgotten might actually be in conflict with something we value, we need to ask how important this value is, how substantial the conflict is, and whether it can be resolved. Given the goal of brevity, all that I can say at this moment is that Garcia-Murillo and MacInnes have not sufficiently considered fixing the application of the right to be forgotten and that the worries I raise arguably outweigh the worries they raise about the right to be forgotten. In particular, I believe this to be true about the manipulation-based worries of Garcia-Murillo and MacInnes' proposal. As I say, my worry is not only that they do no resolve the manipulation worry, but that their proposal increases it (see fn. 1, fn. 9, and the end of the previous section). In fact, I believe that this problem alone is reason to reject their proposal.

Funding

This work is part of the research programme Ethics of Socially Disruptive Technologies, which is funded through the Gravitation programme of the Dutch Ministry of Education, Culture, and Science and the Netherlands Organization for Scientific Research (NWO grant number 024.004.031).

Acknowledgement

I want to thank Joel Andersson, for comments prior to submissions. I also want to thank three reviewers for this journals.

References

Bicchieri, C. (2017). *Norms in the wild: How to diagnose, measure, and change social norms*. New York, NY: Oxford University Press.
Brin, D. (1998). *The transparent society: Will technology force us to choose between privacy and freedom? Reading, mass*. Addison-Wesley.
Bursell, M., Bygren, M., & Gählerab, M. (2021). Does employer discrimination contribute to the subordinate labor market inclusion of individuals of a foreign background? *Social Science Research, 98*, Article 102582. https://doi.org/10.1016/j.ssresearch.2021.102582
Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science, 9*(3–4). https://doi.org/10.1561/0400000042
Foucault, M. (1995). *Discipline and punish: The birth of the prison* (2nd Vintage books ed). New York: Vintage Books [1977].
Fried, C. (1970). *An anatomy of values: Problems of personal and social choice*. Cambridge, Mass.
Garcia-Murillo, M., & MacInnes, I. (2018). Così fan tutte: A better approach than the right to be forgotten. *Telecommunications Policy, 42*(3), 227–240. https://doi.org/10.1016/j.telpol.2017.12.003
Gerstein, R. (1978). Intimacy and privacy. *Ethics, 89*, 76–81. https://doi.org/10.1086/292105
van den Hoven, M. J. (1997). Privacy and the varieties of moral wrong-doing in an information age. *SIGCAS Comput. Soc., 27*(3), 33–37. https://doi.org/10.1145/270858.270868
Johnson, J. L. (1989). Privacy and the judgment of others. *The Journal of Value Inquiry, 23*, 157–168. https://doi.org/10.1007/BF00137284
Lundgren, B. (2020a). The concept of anonymity: What is really at stake? In K. Macnish, & J. Galliot (Eds.), *Big data and democracy* (pp. 201–216). Edinburgh: Edinburgh University Press. https://edinburghuniversitypress.com/pub/media/resources/9781474463522_Chapter_13.pdf.
Lundgren, B. (2020b). A dilemma for privacy as control. *The Journal of Ethics, 24*, 165–175. https://doi.org/10.1007/s10892-019-09316-z
Lundgren, B. (2020c). Against AI-improved personal memory. In J. Haltaufderheide, J. Hovemann, & J. Vollmann (Eds.), *Aging between participation and simulation: Ethical Dimensions of socially assistive technologies in elderly care* (pp. 223–234). Berlin, Boston: De Gruyter. https://doi.org/10.1515/9783110677485-014.
Lundgren, B. (2021a). Ethical machine decisions and the input-selection problem. *Synthese, 199*, 11423–11443. https://doi.org/10.1007/s11229-021-03296-0
Lundgren, B. (2021b). How we can make sense of control-based intuitions for limited access-conceptions of the right to privacy. *Journal of Ethics and Social Philosophy, 20*(3), 383–391. https://doi.org/10.26556/jesp.v20i3.1438
Lundgren, B. (2021c). Confusion and the role of intuitions in the debate on the conception of the right to privacy. *Res Publica, 27*, 669–674. https://doi.org/10.1007/s11158-020-09495-9
Marmor, A. (2015). What Is the Right to Privacy? *Philosophy and Public Affairs, 43*(1), 3–26. https://doi.org/10.1111/papa.12040
McDonnell, M.-H., & King, B. G. (2018). Order in the court: How firm status and reputation shape the outcomes of employment discrimination suits. *American Sociological Review, 83*(1), 61–87. https://doi.org/10.1177/0003122417747289
Mooradian, N. (2009). The importance of privacy revisited. *Ethics and Information Technology, 11*, 163–174. https://doi.org/10.1007/s10676-009-9201-2
Munch, L. A. (2020). The right to privacy, control over self-presentation, and subsequent harm. *Journal of Applied Philosophy, 37*, 141–154. https://doi.org/10.1111/japp.12384
Nissenbaum, H. (2010). *Privacy in context: Technology, policy, and the integrity of social life*. Stanford, Calif.: Stanford Law Books.

[9] In all fairness to Garcia-Murillo and MacInnes, they do think that manipulation is a problem, but as I note in fn. 1, their solution to this problem is very limited and my point here is that their proposal would actually escalate the problem.

Nissenbaum, H. (2011). A contextual approach to privacy online. *Dædalus, the Journal of the American Academy of Arts & Sciences, 140*(4), 32–48. https://doi.org/10.1162/DAED_a_00113
Rachels, J. (1975). Why privacy is important. *Philosophy & Public Affairs, 4*(4), 323–333. https://doi.org/10.2307/2265077
Véliz, C. (2020). *Privacy is power: Why and how you should take back control of your data.* London: Bantam Press.