CENTER FOR
THE GOVERNANCE
OF CHANCE

DATA, PRIVACY, AND THE INDIVIDUAL

PRIVACY, AUTONOMY, AND PERSONALISED TARGETING

KARINA VOLD, JESS WHITTLESTONE
UNIVERSITY OF CAMBRIDGE

NOVEMBER 2019

PRIVACY, AUTONOMY, AND PERSONALISED TARGETING: RETHINKING HOW PERSONAL DATA IS USED

Karina Vold, Jess Whittlestone

Leverhulme Centre for the Future of Intelligence

University of Cambridge

Reference to this paper should be made as follows:
Vold, K., Whittlestone, J. (2019): "Privacy, Autonomy, and Personalised targeting: Rethinking How Personal Data is Used", *Data, Privacy and the Individual*

INTRODUCTION

Technological advances are bringing new light to privacy issues and changing the reasons for why privacy is important. These advances have changed not only the kind of personal data that is available to be collected, but also how that personal data can be used by those who have access to it. We are particularly concerned with how information about personal attributes inferred from collected data (such as online behaviour), can be used to tailor messages and services to specific individuals or groups. This kind of 'personalised targeting' has the potential to influence individuals' perceptions, attitudes, and choices in unprecedented ways. In this paper, we argue that because it is becoming easier for companies to use collected data for influence, threats to privacy are increasingly also threats to personal autonomy—an individual's ability to reflect on and decide freely about their values, actions, and behaviour, and to act on those choices. 1 While increasing attention is directed to the ethics of how personal data is collected, we make the case that a new ethics of privacy needs to also think more rigorously about how personal data may be used, and its potential impact on personal autonomy.

We begin by briefly reviewing existing work on the value of privacy and its link to autonomy, before outlining how recent technological advances are changing this relationship by changing the ways that personal information can be used to influence behaviour. We introduce the idea of 'personalised targeting', and discuss its implications for autonomy, before finally presenting some considerations for determining when this kind of targeting is acceptable and when it is not. Finally, we conclude with some practical implications for thinking about the ethics of how data is used.

WHY VALUE PRIVACY?

Other authors in this project have considered the value of privacy in more detail, so here we just briefly discuss the historical link between privacy and personal autonomy. While a few thinkers have suggested that privacy may have intrinsic value (e.g., Fried 1970; Moor 1990), most discussions focus at least partially on the instrumental value of privacy for protecting other goods such as autonomy, dignity, fairness, reputation, self-development, intimacy, and bodily integrity (for overviews, see Solove 2006, 2008). While we think that all of these links are important, here we focus on autonomy, which we believe has a particularly important connection to privacy today.

The basic intuition behind the idea that privacy and autonomy are connected is as follows. One natural reason why a person might care about threats to their privacy is that

[1] This definition is meant to capture some consensus around the notion of 'autonomy', although there continue to be disagreements amongst political philosophers on how best to characterize the notion. See Bernal (2014), Raz (1986), Rawls (1999), and Macnish (2019) for further discussions.

if others have access to personal information about them, others can use that information to influence them. If such influence is exerted in covert or manipulative ways, this could especially threaten a person's ability to make independent decisions and form independent beliefs or values.

One of the earliest accounts of a 'right to privacy', proposed by Warren and Brandeis (1890), came as a response to the nascent inventions of photography and newspapers at the time, which introduced new threats to an important right they termed 'the right to be let alone'. The protection from scrutiny, interruption, and criticism that solitude provides is important for autonomous decision-making. Even this early defence of privacy provides some link to personal autonomy, if only implicitly. Since then, however, others have argued more explicitly for this connection.

On one view, exercising autonomy requires being able to detach from the influences that come from social and political spheres. Westin (1967), for example, argues that the right to privacy protects individual autonomy from interference and influence by society and the state by carving out a protective space for the individual to reflect and act freely. A similar case has been made by other thinkers: Benn (1984), Johnson (1985), and Kupfer (1987) all argue along similar lines that privacy supports personal autonomy. In particular, Kupfer contends that privacy is essential for 'the development of an autonomous self'. For instance, the total loss of privacy in most prisons causes inmates to suffer from a shrunken individual 'self-concept' and, as a result, to suffer from diminished autonomy (Kupfer 1987: 83).2 To be free to self-determine and maintain a robust 'self-concept', therefore, individuals need to be able to retain some control over what information about them is accessible to other people.

While the link between privacy and autonomy has been acknowledged in some academic literature on internet privacy (e.g., Bernal 2014), this relationship has rarely been emphasised in more contemporary policy discussions around digital privacy. The most recent report from the UK Information Commissioner's Office (ICO) on personal information and political influence, for example, never once mentions autonomy (ICO 2018). In this paper we argue that as technological advances make it easier than ever for personal information to be used to influence people in manipulative ways, the link between privacy and autonomy is stronger than ever and deserving of more attention. We begin by outlining in more detail some of the key relevant technological advances around the collection, access, and use of personal data.

[2] Jeremy Bentham (1995) makes a similar case with his example of the hypothetical 'Panopticon', a prison with extreme surveillance that, so he argued, had the effect of reducing individual autonomy.

KEY TECHNOLOGICAL ADVANCES AROUND PERSONAL DATA

While the values of privacy and autonomy have always been interconnected, technological advances have made this link stronger than ever, as we are seeing changes in: (i) the type and amount of personal data that can be collected; (ii) who holds and has access to that data; and, perhaps most importantly, (iii) how that data can be used by those who can access it. In this section we review some of the key technological advances in recent years that are relevant to privacy and autonomy.

(i) Data collection: It is now possible for companies to collect vastly more data about people's lives than ever. A wealth of data about our online behaviour can be accessed (and is sometimes owned) by tech companies. Such data includes web browsing logs, search engine activity, and our social media networks and activity. Much of this data may not be legally considered 'personal data' or at least is not given the strongest legal protections,3 though with machine learning methods it is possible to draw highly personal inferences from it. As such we will use the term 'personally relevant data' to refer to data that is either legally considered to be personal data under the European Union's General Data Protection Regulation (GDPR) or data that could be used (with the right techniques) to draw inferences about a particular human user's personal attributes. For example, age and gender can be predicted with relatively high accuracy from people's web browsing logs (Hu et al. 2007). Facebook 'likes' can be used to predict gender, sexual orientation, religious beliefs, and ethnicity with high accuracy, and to predict personality variables with reasonable reliability (Kosinski et al. 2012). What counts as personally relevant data is therefore meant to be quite broad—it covers all the kinds of trackable online activity mentioned, but it could also include biometric information gathered from wearables, location information from GPS-enabled devices, and even information held by the state such as health care records, income, and tax information. 4

(ii) Data storage and access: Personally relevant data from smartphones and smart devices largely sits in the hands of a few powerful tech companies, such as Google and Facebook. As individuals conduct more of their activities online, from the things they buy to daily interactions with friends and colleagues, these companies are accumulating vast amounts of information about our day-to-day lives. Access to this personal data can also be bought and sold by different private companies, without the knowledge, much less the consent, of those individuals whom the data concerns. This leads to a highly asymmetric situation: tech companies hold an enormous amount of information about their users, information that they may not even know themselves, while at the same time users know

[3] At least not in the EU, under the General Data Protection Regulation (GDPR); see Wachter and Mittelstadt (Forthcoming).

[4] We use this definition, rather than the definition of personal data under the GDPR or other legislation, in recognition that what is legally considered to be 'personal data' can change as either the law or technology changes (for example, advances in machine learning enabling us to draw increasingly personal inferences from data that was not previously considered to be 'personal')

very little about these companies and how they work. This asymmetry gives large tech companies substantial power over individuals. While asymmetric power relations have always existed between individuals and their governments, for example, the power that tech companies have acquired through their access to data is a cause for concern. Though not without their own problems, democratic governments are at least elected by citizens, meaning they are subject to more public accountability than private companies. The high-level objectives of our government are also typically more aligned with public interests than the objectives of a private company might be.

(iii) Data use: A large part of what makes this data so valuable to companies is that it can be used to 'target' the delivery of various products, services, and messages to specific users and demographics, allowing companies to better achieve whatever their business aims are. Targeting is the process of selecting a specific group of people and tailoring messages or services based on the characteristics of that group, as a strategy for more effective campaigns. For example, Facebook adverts and Amazon product suggestions are tailored to the user using an algorithm that takes into account various sources of information, such as past buying behaviour or demographic group. While the question of what will be effective for different types of people has long been integral to any product development or marketing strategy, technological advances are changing the kind of targeting that is possible, making it easier to target groups and individuals.

The process of selecting a specified group is essential for effective targeting. Different products and different types of messaging will appeal to different groups of people—for example, middle-aged women are more likely than teenagers or most men to be interested in anti-ageing cream, and a person's political persuasion may have a significant impact on what kinds of news sources they read. The more information is available about a person's demographic and interests, the easier it is to target messages and services with personalised content to be maximally effective. Through the collection of increasingly personal data described in (i), companies can learn more about their consumers than ever before, making their targeting methods more narrow and personalised. In the future, we may see companies use increasingly sophisticated techniques to make inferences that can inform their targeting efforts—inferring personal characteristics such as personality traits, interests, or health conditions, for example, and using those characteristics to predict or influence future behaviour.

As well as making targeting efforts more personalised, technological advances allow targeting to be made both more pervasive and more covert. Because most people carry smartphones and access websites individually (as opposed to, say, how families used to consume television together), companies have unprecedented access to data on an individual level. Individualised data allows companies to customize the means of their interventions: they can pay for ads on the websites a specific user visits, on the social media feeds they frequent the most, and even geo-target ads based on their current physical location. This means that individuals may be subject to some form of targeting from companies in a large proportion of their daily lives, while being completely unaware that this is happening most of the time.

PERSONALISED TARGETING AND AUTONOMY

(i) What is personalised targeting?

We refer to the entire process of using personally relevant data to customize both the content and the means of interventions as 'personalised targeting' (others have used the label 'behavioral micro-targeting', e.g,. Ward 2018). In this section we outline in more detail what we mean by personalised targeting, how it can be used, and why it is a threat to autonomy.

The ability to use data to target individuals in different ways distinguishes personalised targeting from the kind of broad targeting of populations that has been used for years, such as running advertisements for children's toys during cartoon shows. Modern personalised targeting is also unique in its potential to use and adapt to real-time data. As mentioned, companies can already geo-target ads based on a person's current location, and in the near future, it may even be possible to assess mood and emotions in real-time from users' online communication (using machine learning to predict mood from language use) and heart-rate monitoring on wearable devices.5

How effective is this kind of personalised targeting at actually influencing individual behaviour? The evidence is currently mixed, depending on the type of data being used and the aims of targeting. Some evidence suggests that consumer behaviour is more effectively influenced when campaigns are tailored to individuals' personal characteristics (Matz et al. 2017; Goldfarb 2014; Noar et al. 2007), but no direct evidence exists yet that tailoring based on features such as personality type can be effective in changing users' political attitudes or voting behaviour (Resnick 2018). However, it is clear that companies, and even governments, are attempting to influence behaviour in increasingly personalised ways, and as data availability and modelling techniques improve, it is reasonable to expect that personalised targeting will become more effective.

(ii) Why personalised targeting matters for autonomy

The increased use of personalised targeting by companies to deliver messages and services is not all bad, of course, and others have emphasised the potential benefits for both companies and consumers: how online stores can be better designed to fit the user's inferred profile, and how marketing and product recommendations could be improved

[5] A particularly relevant field of study is affective computing which, among other things, aims to develop ways to recognize and influence human emotions through digital devices and data. For more detail see Calvo et al. 2015.

with a better understanding of individuals' attributes and preferences (Kosinski et al. 2012). However, the recent scandal involving Cambridge Analytica, who allegedly used personality traits inferred from Facebook 'Likes' to target political campaigns (Cadwalladr and Graham-Harrison 2018), demonstrates the potential for personalised targeting to be used to influence behaviour in ways harmful to individuals and society.6

Efforts by companies and governments to shape the beliefs and behaviour of the general public are certainly not a new phenomenon. As mentioned before, marketers have long been using the practice of 'segmenting' audiences, tailoring campaigns to make them maximally persuasive to different groups or demographics. In recent years, governments and companies have also begun using insights from behavioural science to more effectively influence citizen and customer behaviour. Based on 'nudge theory', these relatively new methods of influencing behaviour rely on the idea that decision-making is easily influenced by environmental factors, or changes in our 'choice architecture' (Thaler and Sunstein 2009). 'Nudges' attempt to trigger automatic cognitive processes (rather than reflective or deliberate ones), in order to increase the likelihood that an individual will make a certain choice—making healthier options more salient or easier to reach in a canteen, for example.

What makes personalised targeting a particular threat to autonomy, compared to other attempts to influence or 'nudge' behaviour? We argue that it is because personalised targeting is particularly likely to be manipulative. Here we follow Susser et al. (2018), who define manipulation as any attempt to influence others' behaviour that is hidden and subverts an individual's ability to act on their own reasons. On this definition, manipulation contrasts with attempts to persuade which appeal openly to a person's capacity for conscious deliberation and choice, and therefore are consistent with personal autonomy. The difference between manipulation and persuasion is not necessarily a clear cut one, but rather a spectrum, depending on the extent to which the attempt to influence undermines an individual's ability to act on their own reasons.

Personalised targeting differs from traditional marketing and nudging in a few key ways that make it more likely to be manipulative, and therefore more threatening to autonomy.

First, the kind of personalised targeting that is possible today is becoming increasingly 'hidden' from view. This is partly because attempts to influence behaviour are simply more pervasive in our lives: an increasing number of our choices and behaviours are conducted through interactions with digital technologies. We may be so often subject to attempts to influence our decisions and behaviour that we cannot possibly be alert to each attempt. In addition, because of the information asymmetry between technology companies and their users, in most cases an individual cannot even know what

[6] Despite mixed evidence about how effective their techniques were, many consider the attempt to manipulate voting behaviour itself to be a threat to autonomy, e.g., Ward 2018; Zunger 2018.

information a company has about them, let alone how or when that information is being used.

Second, the highly personalised nature of modern targeting also makes it a bigger threat to autonomy. Susser et al. (2018) suggest that while mass manipulation is certainly possible, we ought to worry more about manipulation the more targeted it is—in part simply because the more personalised an attempt at influence is, the more likely it is to be effective. But an additional reason why personalisation may be especially worrying is that personalised attempts to persuade are more likely to be deceptive, which Goodin (1980) suggests is key to manipulation. If a company is attempting to persuade a large, diverse group of people to vote for a certain political candidate, they will to have to use a strategy that will appeal widely, highlighting many of the candidate's different policies and strengths. But if the objective is to appeal to a specific individual, and the company has information on that person's characteristics and interests, a more effective approach might be to just emphasise those policies that are most likely to appeal to that person. If the targeted person is not made aware that this information has been personalised to them, they may assume they are seeing a much more representative picture of the candidate than they in fact are. In an extreme scenario, a candidate could even deliberately mislead voters by presenting herself as in favour of gun control to one group of people, for example, and as against it to another group. As personalised targeting increases companies' ability to restrict or misrepresent the information and options available to them, it therefore increases the threat to individuals' ability to decide and act freely i.e. their autonomy.

Finally, personalised targeting methods are largely being used by powerful technology companies whose goals are not necessarily aligned with the well-being of individual users. Even if companies such as Facebook do not intend to influence their users in harmful ways, they may inadvertently do so simply because the company's objective does not align with individual users' well-being. Facebook's aim is to increase users' engagement with the site. By tailoring a user's experience of the site to optimise that aim, the company is automatically trying to influence users to do what Facebook wants rather than allowing them the autonomy to choose what is best for themselves (which may well be to stop scrolling through their newsfeed.) Powerful technology companies such as Google and Facebook also have the power to shape more than just people's purchases, but also their most fundamental beliefs—about politics, morality, and the reality of what others around them do and think. This presents a much more severe threat to individual autonomy, as these attitudes are arguably much more central to who one is as a person than consumer decisions.

To sum up: while individuals, companies and governments have always attempted to influence others, modern personalised targeting makes these attempts more covert, widespread and deceptive, and generally less aligned with public interests. All of this brings new urgency to familiar concerns about the ethics of influence and manipulation

(Sunstein, 2016), raising new questions about the use of highly personal data in attempts to influence behaviour. When and how is it acceptable to use information about personal attributes (inferred or otherwise) to target adverts, messages, and interventions at specific groups or individuals in order to influence their attitudes and behaviour? And how do we ensure that individual autonomy is protected as the amount and kinds of personal data that are available increases?

WHEN IS PERSONALISED TARGETING ACCEPTABLE?

This section introduces some of the factors that we suggest are most important when considering what should count as an ethical use of targeting, building on what has been written about the ethics and public acceptability of nudging (Selinger and Whyte 2011; Sunstein 2015, 2016; Petrescu et al. 2016).

Some cases of personalised targeting, such as the targeted political campaigns Cambridge Analytica has been accused of engaging in, have been called out by many as highly unethical, both in the ways described above and in threatening the democratic process (e.g., Ward 2018; Zunger 2018). 7 On the other hand, it does seem possible that personalised targeting can be used for good: imagine a hypothetical medical startup which consensually uses data about individuals' personal medical histories to tailor interventions to be maximally effective.

In some sense, Cambridge Analytica and the hypothetical medical startup are engaged in similar practices: both are using personal information to target the 'service' they provide in order to be more effective. Many cases of targeting, both real and hypothetical, will lie somewhere in between these two ends of the spectrum and raise more nuanced ethical questions. What about a government using inferences about personal attributes to deliver personalised tax reminder messages, increasing the likelihood that people will pay their taxes on time? Or targeted advertising based on segmenting populations by demographic factors such as gender, class, or educational background? How do we draw the line between where targeting seems to yield clear benefits and where it looks like outright manipulation? What factors—such as the kind of data being used, the level of transparency, or the purpose of the targeting—make the difference?

The following are some considerations that seem important in distinguishing between ethical and unethical forms of targeting.

[7] Indeed, there are also important social consequences of personalised targeting, such as threats to democratic processes, fairness, and social equality, but this paper focuses on the negative consequences for the individual rather than for society.

(i) Personalised targeting should be consistent with people's values and interests.

First, the objective of the targeting is clearly important. A company might be trying to change people's preferences or behaviour to suit the company's own objectives—say, to get people to vote for a given political party or increase their use of a social media platform, say. Contrast this with a case in which targeting instead aims to better serve existing preferences, such as providing targeted 'nudges' to someone who has purchased a FitBit to help them walk more. A useful question to ask here is the following: to what extent are the interests of the 'targeter' aligned with those of the person being targeted?

(ii) Personalised targeting should be transparent.

A second important factor is how transparent it is to users that targeting is taking place, and what kinds of information are being used for this targeting. Most people are aware that advertisers are giving them personalised content, but may well assume they are seeing the same political campaigns or news articles as everyone else. In most cases people also have no awareness of what assumptions about them are being used in targeting: whether the adverts or campaigns they are seeing are tailored based on their gender, ethnicity, interests, or location for example. When people are aware that they are seeing personalised content, such as in the case of product advertising, they can (a) take this into account when making decisions based on the content they see (recognising that it may not be fully neutral or objective), and (b) have at least some recourse to challenge or change the way such targeting is being done. Both remedies rely on transparency to help people preserve their autonomy.

(iii) Companies should attempt to seek consent.

Related to the issue of transparency is whether people are given an opportunity to consent to how their data is being used for targeting. Ideally, consent should have been obtained for both (i) what personally relevant data is collected and (ii) how that data is used in personalised targeting. In the case of the medical startup, we can assume that a patient's individual medical history was given to the startup with that person's consent and an understanding of how it would be used. This is starkly different from how Cambridge Analytica collected and used personally relevant data in the form of Facebook 'Likes'. Gaining consent can be challenging in practice, however. As many Europeans have experienced since the implementation of the GDPR in 2018, for example, abundant notifications about the use of cookies can easily become a burden for the user (Schofield 2018; Degeling et al. 2018). More research is needed to develop practical methods of gaining consent, given its importance in preserving users' autonomy. One option, already included in the GDPR, is to require that users be able to easily (e.g., with just a few clicks) opt out of the collection of their data at any time. It would also be desirable for them to be able to withdraw consent for certain uses of their data.

(iv) Personalised targeting should not attempt to restrict information or choices in a way that misrepresents reality.

Personalised targeting can misrepresent reality in two ways: it may restrict what information people are aware of, or it may restrict the choices that are made available to them, each undermining different aspects of autonomy. It may be impossible to present someone with all possible information or choices without overwhelming them, and so focusing on what is most relevant does not necessarily restrict autonomy. One important question in this context is: how would information or options be presented in the absence of personally relevant information? The more personalised targeting departs from this 'default', the more it risks limiting individuals' ability to choose and act freely, and therefore should be held to a higher standard of justification.

(v) Personalised targeting should not make use of certain kinds of personally relevant data.

Some types of personally relevant data are more sensitive than others, and there may be certain information that we want to deem strictly unacceptable to exploit for the purposes of targeting. For instance, targeting individuals from minority ethnic groups with ads for alcohol or tobacco products because those groups have higher consumption rates seems clearly problematic (e.g., Moore et al. 1996). Likewise, using information about vulnerabilities—for example, inferring that someone suffers from depression or anxiety, and using this to target persuasion attempts at those vulnerabilities seems particularly unacceptable. There can be exceptions: if information about vulnerabilities is relevant to identifying those who would most benefit from a certain kind of support, for example. In these cases, ensuring the person can consent is particularly important. This brings us back to the importance of considering to what extent the objective of targeting is aligned with the individual's own interests. But to the extent that personally relevant data is used to exploit vulnerabilities, the use of this data is particularly problematic. Of course, determining what counts as a vulnerability in practice will be challenging, and much more work is needed to draw meaningful and useful lines here.

There are a number of additional questions we might ask about what kinds of attributes it is acceptable to use for targeting. Is narrower targeting—that is, on the level of individual characteristics rather than group attributes—more ethically problematic? Is it more acceptable to use attributes that are relevant to the purpose of targeting? Are there some types of personally-relevant information that should always be off-limits for the purposes of targeting?

IMPLICATIONS AND CONCLUSION

This paper contends that now, more than ever, forsaking our privacy threatens individual autonomy. Technological advances in how personally relevant data can be collected and accessed, enable companies and other organisations to influence individuals behaviour in highly targeted ways. It is this process of personalised targeting that raises new threats to autonomy, especially as it is deployed by technology companies whose goals are not necessarily aligned with individuals' well-being.

What does all of this mean for how society thinks about privacy? The contention of this paper is that there is now a need for the ethics of privacy to reevaluate how personal data can be used to undermine individual autonomy, and to take this into account in determining how the collection and use of personally relevant data ought to be governed. As Wachter and Mittelstadt (Forthcoming) point out, for example, some of the information that can be inferred from data about online behaviour can pertain to highly sensitive attributes, but is not currently protected by the law to any serious degree. Especially given how these inferences about personal attributes might be used for influence, there may be a strong case for changing the governance of such information.

In this paper, we have outlined five relevant factors that an ethics of personalised targeting must consider in order to better protect people's autonomy: (1) to what extent the objectives of the targeting effort align with the individual's own values and interests, (2) how transparent the targeting effort is, (3) whether the individual has given consent for their data to be collected and used for the purposes of targeting, (4) whether the form of targeting makes a fair attempt to accurately represent reality, and, finally, (5) whether certain sensitive, vulnerable information about an individual has been used.

The last two considerations are especially distinctive to new forms of personalised targeting,8 and so particularly warrant further exploration, perhaps using different case studies to delineate key ethical considerations. In particular, we need better answers to the following questions: (a) what counts as making a fair attempt to 'accurately represent reality' in personalised targeting?, and (b) what kinds of sensitive information should be considered off-limits for the purposes of targeting, and in what contexts?

We urge both policymakers and researchers to take the ethics of personalised targeting seriously in discussions of privacy, and to consider how existing regulation or new governance approaches might be used to protect individual autonomy. Considerations about how data may be used, and the ethical consequences of that use, need to be a part of the larger analysis of what data it is acceptable for companies and others to acquire.

[8] Versions of our first three principles also appear in Sunstein's (2015) work on the ethics of nudging.

REFERENCES

Benn, S. I. (1984). Privacy, freedom, and respect for persons. In F. Schoeman (Ed.), Philosophical dimensions of privacy: An anthology (pp. 223-224). Cambridge, UK: Cambridge University Press.

Bentham, J. (1995). The Panopticon Writings. Ed. Miran Bozovic. London: Verso Books.

Bernal, P. (2014). Internet Privacy Rights: Rights to protect autonomy. Cambridge University Press.

Cadwalladr, C., & Graham-Harrison, E. (2018). How Cambridge Analytics turned Facebook 'likes' into a lucrative political tool. The Guardian. March 17, 2018. https://www.theguardian.com/technology/2018/mar/17/facebook-cambridge-analytica-kogan-data-algorithm. Accessed January 4, 2019.

Calvo, R. A., D'Mello, S. K., Gratch, J. & Kappas, A. (2015). (Eds). Handbook of Affective Computing. Oxford University Press.

Degeling, M., Utz, C. Lentzsch, C., Hosseini, H., Schaub, F., & Holz, T. (2018). We value your privacy…now take some cookies: Measuring the GDPR's impact on web privacy." ArXiv, arXiv:1808.05096

Preprint available online: https://arxiv.org/pdf/1808.05096.pdf. Accessed January 3, 2019.

Fried, C. (1970). An anatomy of values. Cambridge, MA: Harvard University Press.

Goldfarb, A. (2014). What is different about online advertising? Review of Industrial Organization, 44(2), 115–129. https://doi.org/10.1007/s11151-013-9399-3

Goodin, R. E. (1980). Manipulatory politics. New Haven: Yale University Press.

Hu, H.-J. Zeng, H. Li, C. Niu, & Chen, Z., (2007). Demographic prediction based on user's browsing behavior. In Proceedings of the 16th international conference on World Wide Web, 2007 Banff, AB, ACM: pp. 151-160.

Information Commissioner's Office (ICO) of the United Kingdom. (2018). Democracy disrupted: Personal information and political influence. Self-published report. July 11, 2018. https://ico.org.uk/media/action-weve-taken/2259369/democracy-disrupted-110718.pdf. Retrieved on January 4, 2019.

Johnson, D. (1985). Computer ethics. Englewood Cliffs, NJ: Prentice Hall.

Kosinski, M., Kohli, P., Stillwell, D. J., Bachrach, Y., & Graepel, T. (2012). Personality and website choice. ACM Web Science Conference, Evanston, Illinois, USA, 251–254.

Kupfer, J. (1987). Privacy, autonomy, and self-concept. American Philosophical Quarterly 24(1), 81–89.

Matz, S. C., Kosinski, M., Nave, G., & Stillwell, D. J. (2017). Psychological targeting in digital mass persuasion. Proceedings of the National Academy of Sciences 114 (48), 12714–12719. DOI:10.1073/pnas.1710966114

Moor, J. H. (1990). The ethics of privacy protection.Library Trends 39(1–2), 69–82.

Moore, D. J., Williams, J. D., & Qualls, W. J. (1996). Target marketing of tobacco and alcohol-related products to ethnic groups in the United States. Ethnicity and Disease 6, 83–98.

Noar, S. M., Benac, C. N. & Harris, M. S., (2007). Does tailoring matter? A meta-analytic review of tailored print health behavior change interventions. Psychological Bulletin, 133(4), 673–693.

Petrescu, D. C., Hollands, G. J., Couturier, D. L., Ng, Y. L., & Marteau, T. (2016). Public acceptability in the UK and USA of nudging to reduce obesity: The example of reducing sugar-sweetened beverages consumption. PLOS ONE 11(6): e0155995. https://doi.org/10.1371/journal.pone.0155995

Rawls, J. (1999). Collected Papers. S. Freeman (Ed.) Harvard University Press.

Raz, J. (1986). The Morality of Freedom. Oxford University Press.

Resnick, B. (2018). Cambridge Analytica's "psychographic microtargeting": What's bullshit and what's legit. Vox News. March 26, 2018. https://www.vox.com/science-and-health/2018/3/23/17152564/cambridge-analytica-psychographic-microtargeting-what. Retrieved on January 4, 2019.

Schofield, J. (2018). What should I do about all the GDPR pop-ups on websites? The Guardian. July 5, 2018. https://www.theguardian.com/technology/askjack/2018/jul/05/what-should-i-do-about-all-the-gdpr-pop-ups-on-websites. Accessed January 3, 2019.

Selinger, E., & Whyte, K. (2011). Is there a right way to nudge? The practice and ethics of choice architecture. Sociology Compass 5(10), 923–935.

Solove, D. (2006). A taxonomy of privacy. University of Pennsylvania Law Review 154(3), 477–560.

Solove, D. (2008). Understanding privacy.Cambridge, MA: Harvard University Press.

Sunstein, C. R. (2015). The ethics of nudging. Yale Journal on Regulation 32, 413–450.

Sunstein, C. R. (2016). The ethics of influence: Government in the age of behavioural science. New York, NY, US: Cambridge University Press.

Susser, D., Roessler, B., & Nissenbaum, H. (2018). Online Manipulation: Hidden Influences in a Digital World. Available at SSRN 3306006.

Thaler, R. H., & Sunstein, C. R. (2009). Nudge: Improving decisions about health, wealth, and happiness.New York: Penguin Books.

Ward, K. (2018). Social networks, the 2016 US presidential election, and Kantian ethics: Applying the categorical imperative to Cambridge Analytica's behavioral microtargeting, Journal of Media Ethics 33(3), 133–148, DOI: 10.1080/23736992.2018.1477047

Warren, S., & Brandeis, L. (1890). The right to privacy. Harvard Law Review 4, 193–220.

Wachter, S., & Mittelstadt, B. (Forthcoming). A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI. Columbia Business Law Review. Available at SSRN: https://ssrn.com/abstract=3248829

Westin, A. (1967). Privacy and freedom, New York: Atheneum.

Zunger, Y. (2018). Computer science faces an ethics crisis. The Cambridge Analytica scandal proves it. The Boston Globe: Ideas. March 22, 2018. https://www.delaat.net/ecis/Computer-science-faces-an-ethics-crisis-The-Boston-Globe.pdf. Retrieved January 3, 2019.