Five fears about mass predictive personalisation in an age of surveillance capitalism

Yeung, Karen

DOI:
[10.1093/idpl/ipy020](https://doi.org/10.1093/idpl/ipy020)

License:
None: All rights reserved

Document Version
Peer reviewed version

Citation for published version (Harvard):
Yeung, K 2018, 'Five fears about mass predictive personalisation in an age of surveillance capitalism', *International Data Privacy Law*, vol. 8, no. 3, pp. 258-269. <https://doi.org/10.1093/idpl/ipy020>

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

This is a pre-copyedited, author-produced version of an article accepted for publication in *International Data Privacy Law* following peer review. The version of record: Yeung, Five fears about mass predictive personalization in an age of surveillance capitalism, Vol 8, Issue 3, August 2018, 258-269 is available online at: <https://academic.oup.com/idpl/article/8/3/258/5198970>
<https://doi.org/10.1093/idpl/ipy020>

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Five Fears About Mass Predictive Personalisation in an Age of Surveillance Capitalism

International Data Privacy Law
Special issue on consumers and consumer protection

Karen Yeung

Introduction

Contemporary production and consumption practices are currently being transformed in ways that depart radically from those which characterised the pre-digital, networked age. Indeed, many commentators claim that the 'New Industrial Revolution' that is now dawning will provoke changes across every aspect of social life of a magnitude and scale as disruptive and far-reaching as those brought about by the original Industrial Revolution.¹ While there are many factors that have contributed to these changes, there is no doubt that technological innovation has played a critical role, particularly the emergence of the internet, the ability to store digital data on the cloud, and the rise of ubiquitous computing (including the rapid and widespread take up of 'smart' connected devices in contemporary industrialised societies). All of these innovations have supported the emergence of powerful technologies currently referred to as 'artificial intelligence.' It is these technologies that, in retrospect, may come to be regarded as the 21st century equivalent of the late 19th century's steam engine, for it is the application of machine learning algorithms applied to massive volumes of digital data, that is powering the transformations in cultures of consumption and production that are currently unfolding.

One of the most commercial valuable applications of these technologies entails the digital profiling of individuals and groups across a population. These techniques are used to collect data gleaned from tracking the online behaviour of individuals across a population, and subjecting that data to machine learning techniques in order to create detailed profiles capable of generating highly accurate predictions about the behaviours, interests, preferences and traits of individuals and groups. Because these techniques can operate continuously and automatically, updating their operation via feedback loops which enable algorithms to improve their own performance, they have become exceptionally popular and powerful vehicles for accurately sifting, sorting and scoring individuals across a population in ways that retailers and marketers in a pre-networked digital age could barely have imagined.

Yet despite the sophistication of these 21st century techniques, the predominant business model that has emerged for consumer-facing contemporary digital services and with which the New Industrial Revolution is increasingly associated is rather primitive, for it is essentially one of 'barter'. According to this model, users agree to continuous monitoring of their online behaviour and the collection of digital breadcrumbs thereby generated, in return for services – known in contemporary parlance as a 'free' rather than a 'fee' for services revenue model.² The widespread adoption of

Interdisciplinary Professorial Fellow in Law, Ethics and Informatics, School of Law and School of Computer Science, the University of Birmingham, Distinguished Visiting Fellow, Melbourne Law School. I am grateful to Fiona de Londras for comments on earlier drafts. All errors are entirely my own.

¹ dana boyd and Kate Crawford, 'Critical Questions for Big Data' 15 *Information, Communication and Society* (2012): 662-79.

² Jose van Dijck 'Datafication, dataism and dataveillance: Big Data between scientific paradigm and ideology' (2014) 12 *Surveillance & Society* 199-208. By offering services for 'free' this eliminates an important barrier to adoption faced by firms seeking to attract new customers with initially high uncertainty about their

profiling technologies to analyse and predict the preferences and behaviours of individuals and groups, together with this 'free services' business model, has spawned a post-industrial form of capitalist production which Shoshana Zuboff has dubbed 'surveillance capitalism'.³ According to Zuboff, the logic of surveillance capitalism rests on the generation of revenue from data assets acquired by ubiquitous automated operations, driven primarily by Silicon Valley's hyperscale technology firms (and spearheaded by Google) which achieve growth mainly by leveraging automation via global digital platforms.⁴ The resulting pools of data which these socio-technical systems routinely and strategically collect constitute a new asset class which Zuboff dubs 'surveillance assets'.⁵ Investors (so-called 'surveillance capitalists') generate profit from the global networked environment through a default business model where company valuations routinely depend upon capturing human attention ('the market for eyeballs') rather than revenue as a predictor of return on investment. Surveillance capitalism therefore entails the continuous and pervasive collection of personal data by tracking the digital footprints of individuals, then channeling and controlling flows of personal information while converting them (via the use of AI technologies) into flows of profit, all in ways that are highly opaque to their users.⁶

The gradual but rapid replacement of industrial capitalism with surveillance capitalism can be understood as driven by changes in modes of production made possible by advances in networked computational technologies, highlighting the role of technological innovation as a catalyst for social transformation. It was the development of the steam engine that provided the technological driver underpinning the original Industrial Revolution, and which, in turn, enabled the development of manufacturing and production processes at a scale and magnitude that was not previously possible. The steam engine enabled the construction of machinery housed within factories that could produce identical units on a mass basis. In a similar vein, the development of big data analytic techniques and the widespread and increasing use of connected smart devices which are rapidly becoming indispensable in contemporary everyday life has spawned not only a new model of capitalism, but also a new mode of production: that of 'mass personalisation'. Common to the modes the production prevailing in the industrial era and the emerging networked digital era is their capacity to operate at *scale*. One of the defining features of the first Industrial Revolution that can be understood as critical to its socially disruptive effects was its departure from small-scale production that characterised modes of production within rural, agrarian societies which the emergence of factories made possible, precipitating the growth and development of urban centres which enabled workers to live in close proximity to factories and for which human labour was essential to enable successful mass production.

Despite this common capacity to operate at scale (and thus on a 'mass' basis), mass personalisation can be distinguished from mass production in at least five ways: (1) It is primarily concerned with the provision of *services* (although it is increasingly including the personalisation of *goods*)⁷; (2) rather

valuation of the services offered: Anja Lambrecht, 'The Economics of Pricing Services Online' in S. N. Surlauf and L. E. Blume (eds) *The New Palgrave Dictionary of Economics* (Palgrave Macmillan, 2013).

³ Shoshana Zuboff, 'Big Other: Surveillance Capitalism and the Prospects of an Informal Civilization' (2015) 30 *Journal of Information Technology* 75-89.

⁴ Ibid.

⁵ Ibid, 80. See also World Economic Forum, *Personal Data: The Emergence of a New Asset Class* (2011).

⁶ Ibid.

⁷ The Deloitte Consumer Review (2015) *Made to Order: The Rise of Mass Personalisation*. Available at <https://www2.deloitte.com/tr/en/pages/consumer-business/articles/made-to-order-the-rise-of-mass->

than generating identical units of production, services are *personalised* to each user, tailored to fit his or her individual's tastes, interests, preferences, lifestyle and behaviours; (3) service provision operates, by default, on a *predictive* basis due to the application of advanced algorithmic profiling in order to *infer and predict* each user's service preferences with the aim of *anticipating* user needs and 'pushing' personalised services to them accordingly; (4) services are continually and automatically reconfigured in light of the feedback gleaned from monitoring the recipient's response to the service without requiring the recipient to provide active and intentional feedback concerning their interest in the service thereby provided (5) although recipients of the service can be characterised as consumers, in the same way that those consuming factory-produced outputs under industrial capitalism may be understood as consumers, users are also concurrently providing raw materials and labour to the producer, comprised of both the personal data which they generate via their digital interactions and the resulting 'data exhaust', the digital by-products incidentally generated from these interactions, and which can be used by the service provider to train and improve their algorithms.

This move to mass personalisation is likely to signify a major shift in our *modes and culture of consumption*, and it is clearly ripe for serious and wide-ranging academic research reflection. It is beyond the scope of this article, however, to reflect comprehensively on the broader implications of its emergence. Rather, the purpose of this paper is modest and limited. Although service providers (particularly social media platforms) typically portray personalised services as offering users a more 'meaningful' experience, the shift from mass production to mass personalisation generates a number of potential dangers, including risks to our collective values and normative commitments and which are, in my view, both necessary and indispensable if democratic communities are to thrive. Yet these threats have been largely overlooked in public and academic debates notwithstanding growing public anxieties about the adverse implications associated with the so-called 'rise of the machines'. This paper identifies five fears (or worries) that the rise of mass predictive personalisation may portend for these collective values and commitments which can be grouped into two broad sets of concerns. The first three fears are largely concerned about the way in which these techniques may lend themselves to unfair practices, and which therefore directly implicate the values of *fairness and justice* which can ultimately be attributed to the systematic use of digital profiling techniques that apply machine learning algorithms to merged sets of data collected from the digital traces generated from continuously tracking users' on-line behaviour to make calculated predictions about individuals across a population. The remaining two fears coalesce around concerns for *social solidarity and loss of community* that may be associated with the increasing personalisation of services and offerings, which is both fuelling and being fuelled by, an increasingly narcissistic mindset that mass personalisation makes possible. In so doing, my aim is primarily to provoke critical discussion and reflection that will motivate more penetrating academic research and place questions of this kind more firmly onto the policy agenda of law-makers and regulators who are in a position to identify and implement the need for safeguards to maintain and protect core values that are necessary for individuals and communities to flourish in networked digital societies.

A more meaningful experience and an empowered consumer?

For the marketing industry, the emergence of digital technologies that enable the automatic personalisation of offers and services to fit the preferences, interests and tastes of individual consumers is in many ways the fulfilment of its ambitions which emerged during the 1950s. During

[personalisation.html](#) (Accessed 16 August 2018). A. Sadiq (2018) 'The shift from mass production to mass personalisation', 31 May, CapGemini Consulting. Available at <https://www.capgemini.com/consulting-gb/2018/05/the-shift-from-mass-production-to-mass-personalisation/> (accessed 16 August 2018).

that time, the retail industry began to attempt to segment their clients' customer base in order to distinguish between existing and potential customers.⁸ For users and consumers, the predictive personalisation of services may appear to offer considerable benefits. Proponents of contemporary personalisation strategies that utilise data profiling techniques typically invoke the rhetoric of consumer 'empowerment'⁹, claiming that they relieve users from unwanted or irrelevant offers whilst the predictive nature of these strategies offers greater levels of convenience and efficiency. For individuals who rank at the top end of what are, in essence, customer scoring systems¹⁰ and thus deemed to be highly desirable customers, mass personalisation is likely to serve them well. But will the turn to predictive personalisation serve consumers as a whole well? And will even those high-end consumers deemed highly desirable be well-served by mass predictive personalisation in the long run? Affirmative answers to these questions are, in my view, far from certain, given the following dangers which such practices are likely to provoke.

Unfair practices

Fear #1: It expands opportunities for consumer exploitation

Because personalisation practices rely on the digital profiling of individuals by mining digital datasets to infer and predict the tastes, interests, preferences and vulnerabilities of individuals, these practices foster and exacerbate the asymmetry of power between profilers (and their industry clients) and those to whom personalised services are provided, thereby increasing the opportunities for the former to exploit the latter. In particular, it is important to bear in mind that contemporary personalisation practices rely critically on the mass surveillance of individuals across populations on a continuous and highly granular basis, and this will invariably have implications for the individuals and groups within the surveilled population. As the Rathenau Institut commented in its report for the Council of Europe's Parliamentary Assembly:

'many technologies nowadays can operate at a distance, most of us are not even aware of the mass surveillance taking place by state and market actors. This creeping development as a whole, and its impact on human rights and society, has received little attention and there has been scarcely any fundamental political and public debate so far. As a result, human beings are rather defenceless relative to this mass surveillance culture, since there are few opportunities to escape the surveillance activities if one does not want to be measured, analysed or coached (as part of a persuasion strategy).'¹¹

These observations not only draw attention to the pervasive population-wide surveillance that mass personalisation necessarily entails, but highlights both the opacity of these processes and the asymmetry in power between the profilers and the profiled. While the former acquire ever greater and more detailed knowledge of the latter, the latter have no equivalent access to the inner workings and practices of the organisations who profile them. In any relationship, whether commercial or otherwise, those with greater knowledge of the other's tastes, interests, preferences, dislikes and vulnerabilities will be better placed to exploit them. At the same time, by dividing consumers into increasingly smaller segments, this further enhances the power of profilers and

⁸ Nora A Draper and Joseph Turow, 'Audience Constructions, Reputations and Emerging Media Technologies: New Issues of Legal and Social Policy' in R. Brownsword, E. Scotford and K. Yeung (eds) *The Oxford Handbook of Law, Regulation and Technology* (Oxford University Press 2017) 1143-1167.

⁹ Sadiq, *supra* n.7.

¹⁰ Frank Pasquale, *The Black Box Society* (Harvard University Press 2015).

¹¹ Rinie van Est and Joost Gerritsen, *Human Rights in the Robot Age* (Rathenau Institut, 2017) 43.

those on whose behalf they act. Such practices resonate directly with the time-honoured strategy of 'divide and conquer' that has been successfully invoked by ruthless political leaders and others to shore up their own power in pursuit of self-interest. Because personalisation strategies serve to isolate individual consumers from each other, they thereby erode consumers' power to act collectively in ways that might serve their interests as a whole. Moreover, because these profiling practices are highly opaque, dynamic and largely operate hidden from public view, this exacerbates and reinforces the asymmetry of knowledge and power between the producer and consumer, thereby increasing the risk of exploitation by the former over the latter.

Some of the most obvious risks of exploitation arise as we enter the so-called era of 'artificial emotional intelligence.'¹² In 2017, it was reported that tech firms are already developing billboards that can recognize and categorize individuals, then demographically direct personalized messages.¹³ Analysis of the click-through behaviour of individuals can readily identify when individuals are feeling low, more likely to make impulse purchases, or more susceptible to particular kinds of offers, enabling retailers to exploit detailed knowledge inferred from user profiles to microtarget personalised offers in ways that will maximise the opportunities to make a sale. But mass personalisation practices need not prey upon user's emotional vulnerabilities to be classed as exploitative. For example, concrete examples may be found in the use of 'personalised pricing', which data-driven profiling and the rise of digital retailing makes possible. Under industrial capitalism, goods were mass produced and supplied to retailers and then offered for sale in bricks and mortar retail stores and made available to all customers entering the store at a particular time at the same price. In contrast, data-driven profiling now enables goods and services to be offered to potential customers at 'personalised' prices. Because each customer only sees his or her own individualised 'digital shop front', and does not have access to the prices or offers made to others on-line, the prices can be personalised through the use of data-driven profiling in order to identify each individual's estimated maximum 'willingness to pay'. While this strategy may optimise revenue for the retailer, it means that two individuals might be offered exactly the same item at precisely the same time, yet at very different prices.¹⁴

Although personalised pricing might not be regarded as exploitative per se, nonetheless the ability of retailers to profile their customers and tailor prices and offers accordingly can support specific personalisation practices that have a distinctly exploitative hue. For example, economic models of price discrimination suggest that there is one group of consumers who might be routinely made worse off through the use of on-line personalised pricing: those consumers who are algorithmically identified as those likely to have difficulty making good decisions, either due to lack of knowledge, poor digital literacy or due to consumer disengagement and who therefore do not actively shop around for better offers, nor switch to alternative providers who are willing to offer them better

¹² Richard Yonck, *Heart of the Machine: Our Future in a World of Emotional Artificial Intelligence* (Arcade Publishing 2017).

¹³ Richard Yonck, 'Welcome to the Emotion Economy – Where AI Responds to, and Predicts, Your Feelings', 2 March 2017, Fastcompany Newsletter. Available at <https://www.fastcompany.com/3067810/welcome-to-the-emotion-economy-where-ai-responds-to-and-predicts-your-feelings> (accessed 16 August 2018).

¹⁴ Christopher Townley, Eric Morrison and Karen Yeung, 'Big Data and Personalized Price Discrimination in EU Competition Law' (2017) 36 *Yearbook of European Law* 683-748; Akiva A Miller, 'What Do We Worry About When We Worry About Price Discrimination?: The Law and Ethics of Using Personal Information for Pricing' (2014) 19 *Journal of Technology Law & Policy*: 41-95.

deals than their current provider.¹⁵ These so-called ‘sleepers’¹⁶ might fail to shop around due to sheer laziness and apathy, and if so, one might regard them as ‘fair game’ by suppliers who can ‘get away with’ charging them higher prices. On the other hand, many of these consumers are likely to be vulnerable individuals, including those on low incomes, the elderly, the digitally disempowered (including those with no ready access to the internet), and the poorly educated. The digital footprint of these individuals may superficially suggest that they have a low price-elasticity of demand (meaning that they are not very responsive to, nor care much about, price changes) because they fail to search out the best deal or switch providers when prices rise. Accordingly, these individuals are more likely to be offered higher prices than well informed, digitally engaged and financially savvy consumers who regularly switch prices, and at a higher price than those that would have prevailed under a uniform pricing regime in which the less savvy consumers are effectively ‘protected’ by well-informed consumers who shop around for the best offers.¹⁷ Consumers who fall into this class may find themselves being offered inferior deals to those offered to others in ways that may appropriately be regarded as exploitative, and yet are likely to be completely unaware that they are being subject to inferior treatment.

Fear #2: It enables subtle but powerful manipulation of individuals at scale

Not only is mass predictive personalisation likely to increase opportunities for those who employ such practices to exploit customers, but it also provides digital service providers with powerful tools of manipulation, which can be used to shape the decisions and behaviours of users in subtle yet powerful ways¹⁸. The recent Facebook/Cambridge Analytica scandal, in which it is alleged that data unlawfully harvested from the Facebook profiles of millions of users was utilised for political micro-targeting in ways that may have perverted the outcome of the US 2016 elections and the Brexit 2016 referendum, reveals not only how readily mass personalisation techniques can be exploited and abused, but also how serious and damaging their consequences might be for the health and integrity of democratic political orders. But even in the consumer context, their manipulative potential becomes apparent by contrasting the characteristics of data-driven predictive personalised services with traditional service personalisation, epitomised by the bespoke suit produced by the tailors of London’s famous Savile Row. Handcrafted human personalisation typically entails the client first identifying a suitable craftsman, reviewing the craftsman’s credentials and qualifications to assess his or her suitability to undertake the task, taking into account, for example, the craftsman’s skills, qualifications, capacity and reliability and so forth. Once identified, the client then requests and invites the craftsman to provide a particular service according to a set of requirements that the client explicitly identifies and stipulates. Those requirements might reflect the client’s conventional patterns of taste, preference and needs, but they might not: indeed, they might depart entirely from her previous tastes and have nothing to do with population-wide trends, or the bespoke service or product might be intended for a person other than the person commissioning the service. In contrast, the typical aim of the algorithmic decision-making systems that rely on machine profiling is

¹⁵ Townley et al, *ibid.*, 37.

¹⁶ Ariel Ezrachi and Maurice E Stucke, *Virtual Competition* (Harvard University Press 2016) 114.

¹⁷ Townley et al, *supra* n.14, 37-38.

¹⁸ Concerns about the use of ‘persuasive technologies’ has been a central theme in the recent observations of the Parliamentary Assembly of the Council of Europe. See for example van Est and Gerritsen, *supra* n.11 and Parliamentary Assembly of Council of Europe, *Report on Technological Convergence, artificial intelligence and human rights* (2017) by Rapporteur Jean-Yves Le Deaut available at <http://assembly.coe.int/nw/xml/XRef/Xref-DocDetails-en.asp?FileID=23531&lang=en> (accessed 14 September 2018).

pre-emptively to *infer* the preferences, behaviours, and lifestyles of individuals across a population in order to tailor informational services (whether video content, search engine results, retail products available for purchase, etc) to fit the preferences and interests which the service provider has inferred through the use of behavioural profiling to generate computational predictions concerning those services that are calculated as *likely* to be of interest to individuals users on *pre-emptive* basis, not only *before* the individual has requested the service, but often *without* the relevant individual requesting such a service at all.

This comparison highlights crucial differences between traditional versus automated, algorithmic forms of service personalisation. In the case of traditional, handcrafted service provision, the client explicitly communicates her needs and requirements about the desired service, while in the case of algorithmic personalisation, the service is automatically configured according to the preferences and interests which the service provider has *inferred* about the individual and offered *pre-emptively* without any express request for service by the individual. Because the individual has not explicitly stated her preferences and interests about the service in question (indeed, she may not want the service at all), the potential to utilise predictive personalisation techniques to pursue ends that may not be in the interests of the customer become more apparent once we attend to the underlying aim of these algorithmic systems. What is it, precisely, that these systems seek to optimise, and who has the power to specify that overarching goal? From the perspective of the system owner, the overarching aim of these systems is to channel user's behaviour and decisions in the direction preferred by the system designer or 'choice architect', and hence the system is intentionally configured to optimise whatever variables will generate maximum commercial returns to its owner.¹⁹ Accordingly there is no guarantee that these objectives will align with the longer-term interests and welfare of users whose decisions and behaviours these systems are aimed at influencing.²⁰

It might be argued that these concerns are misplaced, given that individuals are free to choose whether or not to consume them, just as individuals have always been free to reject or ignore so-called 'invitations to treat' offered to them. At the same time, the whole point and purpose of marketing is to present attractive offers that are likely to be of interest to customers with the ultimate aim of promoting the interests of the retailer. In other words, it is inherent in the nature of capitalism generally that retailers will seek to influence the behaviour of customers in order to benefit the financial interests of sellers and producers, and it has long been recognised that the design and placement of products in-store can have significant effects on the purchasing behaviour of customers in-store. But despite these similarities between marketing practices in the pre-digital age and those involved in mass predictive personalisation, it is important to recognise the powerful, subtle and typically subliminal way in which these data-driven, predictive profiling systems operate. These employ particular kinds of 'nudging' techniques²¹ which intentionally seek to exploit the systematic tendency of individuals to rely on cognitive heuristics or mental short-cuts in making decisions, rather than arriving at them through conscious, reflective deliberation. It is their intentional exploitation of the cognitive weaknesses of individuals with the aim of influencing them to behave ways desired by the choice architect which underpins their deceptive qualities.²²

¹⁹ Karen Yeung, "'Hypernudge': Big Data as a mode of regulation by design' (2017) 20 *Information, Communication & Society* 118-136.

²⁰ James Grimmelman, 'The Platform is the Message' (2018) *Georgetown Law Technology Review* 2: 217-233.

²¹ Richard Thaler and Cass Sunstein, *Nudge* (Penguin Books 2008).

²² Karen Yeung 'Nudge as Fudge' (2012) 75 *Modern Law Review* 122-148.

Not only are nudging techniques problematic because they can be understood as manipulative and lacking in transparency,²³ but when used for the purposes of mass predictive personalisation, their manipulative power is enhanced. In particular, networked data-driven profiling techniques make it possible for each user's choice architecture to be continuously reconfigured in real time in three directions: firstly, through continuous refinement of the user's choice environment in responses to changes in the user's behaviour and environment that have been identified by the system designer as relevant to the user's decision-making through algorithmic analysis of the user's constantly expanding data profile; secondly, by feeding back data to the system designer, which can itself be collected, stored and repurposed for other algorithmic applications; and thirdly, by continuous monitoring and refinement of the individual's choice environment in light of population-wide trends that is made possible through population-wide surveillance upon which mass personalisation techniques rely. Data-driven nudging techniques are therefore particularly nimble, unobtrusive and highly potent, all in ways that qualitatively enhance their manipulative power and which I have referred to elsewhere as 'hypernudging'²⁴ and which Lanzig argues can be understood as threatening both informational and decisional privacy.²⁵

Fear #3 It systematically marginalises and excludes 'low value' individuals

The consumer empowerment rhetoric employed by advocates of data-driven personalisation highlight the capacity of these services to filter out and remove irrelevant services and promotional offers from the informational choice environment of users. While this may be beneficial to users, it nonetheless conceals the underlying processes which such personalisation processes entail: that it is the users themselves who are being sifted, sorted, and scored.²⁶ As Rieder has observed, the collapse of traditional social structures of class, religion and lineage and associated heterogeneity in individual lifestyles and preferences creates acute challenges for marketers and retailers wishing to engage in customer segmentation.²⁷ Accordingly, one of the great attractions of data-driven profiling techniques is their capacity to 'make the social legible again' by 'reinstalling mastery over societies that continuously diversity, creating differentiation that no longer conform to traditional groupings and categorisation' and which can be understood as the 'raison d'être' of contemporary data analytics.²⁸ Although those individuals who score highly in these algorithmic rankings are likely to benefit in the form of generous and attractive offers and opportunities, those who score poorly, and are thus deemed poor prospects for marketers and retailers alike, are likely to be disadvantaged and disempowered by the turn to mass personalisation. As Gandy has observed:

²³ *Ibid.*

²⁴ Karen Yeung, "'Hypernudge': Big Data as a mode of regulation by design' (2017) 20 *Information, Communication & Society* 118-136.

²⁵ Marjolein Lanzing, "'Strongly Recommended" Revisiting Decisional Privacy to Judge Hypernudging in Self-Tracking Technologies' (2018) *Philosophy and Technology* doi.org/10.1007.

²⁶ Draper and Turow, *supra* n.8.

²⁷ Bernhard Rieder, 'On the Diversity of the Accountability Problem: Machine Learning and Knowing Capitalism' (2015) 2 *Digital Culture & Society* 39-54.

²⁸ Rieder, *ibid.* citing Mark Andrejevic, *Infoglut: How Too Much Information Is Changing the Way We Think and Know* (Taylor & Francis, 2013) 11.

Part of the difficulty with the collection and use of personal information to support the marketing and sale of a variety of goods and services rests in the fact that personal information is not only used to *include* individuals within the marketing scan, but may also be used to *exclude* them from other life choices linked to employment, insurance, housing, education and credit.²⁹

Almost a quarter of a century ago, Gandy had already warned of the dangers of what he called the 'panoptic sort', through which individuals are monitored, identified and classified for the purposes of enabling retailers and advertisers to direct their efforts towards groups they believe represent good investment while ignoring those who are deemed unlikely to yield a profit.³⁰ From the perspective of those individuals who are systematically excluded from a wide range of opportunities and facilities, many of which have very significant effects on their individual welfare, mass personalisation threatens to diminish their personal autonomy. According to legal philosopher Joseph Raz, autonomy not only requires the capacity to plot one's own life, but it also requires that individuals have a range of acceptable options from which to do so.³¹ Thus, by foreclosing options that would otherwise be available to the individual in the absence of mass personalisation, data-driven profiling techniques can be understood as narrowing range of worthwhile opportunities available to them and this, in turn, curtails their autonomy.³² In societies committed to individual freedom, the autonomy-diminishing impact of mass personalisation should be a real cause for concern.

One of the most problematic characteristics of the data-driven profiling techniques upon which mass personalisation relies is its opacity, arising from both the highly complex nature of these socio-technical systems, their protection as trade secrets which often applies to algorithms developed by commercial providers and through which individuals are sorted and scored, and the way in which they operate automatically, seamlessly and unobtrusively integrate into users' daily routines. This makes it practically impossible for individuals to understand why particular services offered to them at a particular time and on particular terms, or to identify and understand how and why they are being profiled in a certain way and with what consequences for the resulting personalised offers and services they receive. As a result, the opportunity for individuals to challenge or otherwise contest the way that they have been profiled, sorted and scored, particularly when those profiles have been constructed on the basis of erroneous or otherwise inaccurate data, are extremely limited.³³ As Rieder has observed, algorithmic tools offer an 'aura of objectivity, rationality and legitimacy' that is derived from their empirical underpinnings.³⁴ And in cases of systematic exclusion from opportunities, these may be virtually impossible for the excluded individuals to detect. Take, for example, the study by Carnegie Mellon researchers which found that male users were shown high

²⁹ Oscar H Gandy, 'Coming to Terms with the Panoptic Sort' in D Lyon and E Zureik (eds) *Computers, Surveillance and Privacy* (University of Minnesota Press 1996) 132.

³⁰ Oscar H Gandy, *The panoptic sort: a political economy of personal information* (Westview 1993).

³¹ Joseph Raz, *The Morality of Freedom* (Oxford University Press 1986) 398.

³² *Ibid.*

³³ For disturbing examples of the effects on individuals of inaccurate algorithmic profiling, see Bill Davidow, 'Welcome to Algorithmic Prison - The use of Big Data to profile citizens is subtly, silently constraining freedom,' *The Atlantic*, 20 February 2014 and Cathy O'Neil, *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy* (Allen Lane 2016).

³⁴ Bernhard Rieder, 'On the Diversity of the Accountability Problem: Machine Learning and Knowing Capitalism' (2015) 2 *Digital Culture & Society* 39-54.

paying job ads six times more often than their female counterparts, based on automated algorithmic assessment that concluded that women were 'not interested' in high paying jobs because they had historically not been employed in high paying occupations.³⁵ How could female users possibly detect that they these ads were being systematically excluded from being shown to them? Although the GDPR introduces a number of rights which might assist data subjects seeking information of this kind, many of these rights only apply to algorithmic decision-making systems that have 'significant effects' – it remains unknown whether promotional offerings and advertisements would be interpreted to fall within the scope of this term.³⁶ For Turow, these profiles (which he calls 'reputations') play a central role in determining the promotional content (offers) that a person encounters online and these can be hard to shift' (Turow 2011), commenting:

'That marketers and media producers employ these activities hidden from the vast majority of viewers is deeply problematic – for the reputations of individuals may bear little on an individual's self-perception, while viewers have little opportunities to correct misconceptions.' (Draper & Turow: 1161)

Distributive justice, community and social solidarity

Fear #4: It perpetuates structural inequalities and exacerbates distributive injustice

The aggregate and cumulative effect of mass personalisation over time is likely to contribute to and exacerbate social inequality and distributive injustice. By providing marketers and retailers with the technological capacity to segment consuming publics into distinct groups, based on their relative value and profitability to the retailer, these technologies enable sellers to engage in a commercially rational form of social sorting, seeking to cultivate and attract the choicest customers (the so-called 'strong market' in economic parlance) and exclude low- value customers (the 'weak market'). Akiva Miller argues that if a considerable part of the information used as the basis for personalisation pertains to persistent qualities, such as the neighbourhood in which consumers live, their income level and maximum level of educational attainment, and hence is likely to result in some individuals routinely suffering the same poor quality offers and treatment across different sellers. Commenting on the use of personalised pricing, Miller fears that, over time, these practices threatens to create a market divide between a class of consumers who receive lavish personal attention and preferential treatment, including offers for the best products, services and prices, and a class of consumers algorithmically assessed as 'low value' who are systematically ignored. In other words, Miller fears that retailers will, in essence, rationally seek to 'cherry pick' the rich, and 'lemon drop' the poor.

In many ways, mass personalisation can be understood as pursuing the logic of market segmentation to its ultimate destination, in which each individual user is reduced to a unique market. The overarching logic of market segmentation is to identify and reinforce differences between consumers as the basis for identifying the services deemed (from the perspective of the service provider, rather than the user) most suitable to their inferred needs, preferences, traits and interests. While the appeal of customer segmentation strategies have understandable appeal to service providers (assuming that the cost of pursuing these strategies does not outweigh their value to the provider), segmentation in the on-line world can produce and reinforce societal inequalities because of their inevitable exclusionary effects in ways that may be much harder to detect and combat in comparison to the pursuit of segmentation strategies in traditional 'bricks and mortar'

³⁵ Amit Datta, Michael Car Tschantz, and Anupam Datta, 'Automated Experiments on Ad Privacy Settings' (2015) 1 *Proceedings on Privacy Enhancing Technologies* 92-112.

³⁶ Lillian Edwards and Michael Veale "'Slave to the Algorithm'? Why a 'Right to an Explanation' is Probably Not the Remedy You are Looking For' (2017) 16 *Duke Law & Technology Review* 18-84.

retail environments. In the off-line world, market segmentation does not necessitate exclusion of consumers from one market from accessing the services offered in another market on the same terms and conditions available there. For example, supermarkets like Waitrose and Wholefoods Market seek to target and appeal to high-end, high income customers by stocking a greater range of fresh organic produce, luxury and imported groceries, gourmet ready meals and an extensive range of wines, while others such as Aldi and Lidl might seek to target low-end customers, offering a more limited product range across a set of basic groceries, stocked in higher volumes and offered at lower prices. Although, in the off-line world, a customer who routinely shop at Aldi may nonetheless shop at Waitrose every now and again, simply by entering the store and availing himself of the full range of products available in-store at the same prices and on precisely the same terms as those offered to those customers who routinely shop there (leaving aside any loyalty discount that the latter might receive). In other words, the segmentation of customers does not prevent customers from one segment to take advantage of offers targeted at a quite different customer segment, simply by shopping in-store. In the on-line environment, however, this is no longer possible. Although low-end consumers are free to browse the websites aimed at affluent consumers, they have no way of knowing whether the prices at which good and services offered to them are the same as those offered to others, because they only see their own 'personalised' version of the retailer's site. This not only partitions consumers into discrete segments, but also segregates and excludes them from the offers made available to other consumers, and of which they are likely to be completely unaware. Accordingly, the net cumulative effect of mass personalisation at scale is to perpetuate and reinforce existing forms of systematic societal discrimination, yet these inherently discriminatory and exclusionary logics and effects may proceed unnoticed, yet which the rhetoric of personalisation serves to conceal and obscure.³⁷

Fear #5: It fuels a culture of narcissism, prioritising economic morality over social equality thus eroding solidarity and community

My fears about the effect of data-driven mass personalisation techniques to exacerbate social inequality and collective justice point to deeper misgivings about what the cumulative and aggregate effect of these practices may portend for us, both in terms of our moral character as individuals, and our collective culture as a moral and political community. While the tech and digital marketing industries portray personalisation strategies as a boon for users, these are achieved by seeking pre-emptively to cater to the individual's predicted needs, desires, interests and tastes, through reliance on algorithmic sorting, scoring and ranking. Mass personalisation is therefore likely to create and exacerbate distributive injustice (see #4) by furthering and fostering processes of *social fragmentation* through which individuals are classified, categorised and mathematically sorted in ways that reinforce *differences* between individuals and groups across society. Although it is impossible to prove that the increasing polarisation and social segmentation that is evident in results of political elections across many industrialised and ostensibly democratic states in recent years as extremist political parties have gained ground can be attributed to the increasing use of data-driven profiling to personalise media content pushed to users at scale, it is far from implausible to think that these techniques have significantly contributed to it. The socially corrosive effects of this fragmentation are evident in the dangers associated with political manipulation and microtargeting that was alleged utilised at scale in recent presidential election and referendum campaigns in the US and UK respectively.

Yet it is not merely the effect on political processes that forms the focus of my fifth and final fear. Rather, one of my gravest concerns about mass personalisation are the elusive yet highly consequential effects of these practices on both our character and mindset as individuals and on the

³⁷ Rieder, *supra* n.34.

collective moral culture of our societies. Population-wide and pervasive data-driven personalisation prioritise values of economic morality over values associated with community, particularly the value of equality, in ways that may weaken and marginalise our commitment to equality as both a matter of principle and social practice. My fear is that, deployed at scale, the way in which these data-driven personalisation techniques are being applied in contemporary societies to maximise the economic 'value proposition' for the organisations and institutions that utilise them is likely to have two significant and troubling consequences. First, it fosters a cult of the individual that signifies a shift from a culture of capitalism that can be understood as moving beyond a cultural of material consumerism to one of narcissism. Secondly, mass personalisation may over time corrode social solidarity and so loosen our social bonds that it could threaten the very nature of our collective character as a moral and political community. Each of these two concerns are briefly discussed in turn, beginning first with the rise of narcissism.

In one respect, the emergence of mass personalisation can be understood as a straightforward and unexceptional manifestation of the logic of capitalism through which resources and opportunities follow the logic of market forces. Yet it cannot be denied that the widespread practices of the marketing industry have significantly affected our political and social culture, prompting economic sociologists and political economists to examine the broader social and political implications of what is sometimes referred to as 'varieties of capitalism'. In many ways, the introduction of modern marketing techniques from the 1950s onwards, served to inform a logic of industrial capitalism associated with mass production, consumption and advertising. As Draper & Turow observe, in the 1950s, there were already moves towards 'market segmentation', marking a shift in the logic of industrial capitalism away from mass production towards mass customisation, which – by the 20th century, had become entrenched³⁸. At the same time, the logic of neoliberal policies resting on an unshakable belief in laissez-faire free market capitalism pursued by the Thatcher and Reagan administrations on both sides of the Atlantic from the early 1980s onwards, is widely seen as accelerating and deepening the rise of 'consumerism' – in which the increasing consumption of goods seen as economically desirable, and in which individuals are increasingly preoccupied by, and inclined towards, the buying of consumer goods such that material wealth and possessions become an important indicator of self-worth and social status.³⁹ My fear is that because predictive personalisation strategies that surveillance capitalism fuels are all oriented towards fulfilling the idiosyncratic (albeit inferred) preferences, tastes, desires and inclinations of each individual, they may serve to reinforce and legitimate each individual's belief in the central importance of their own personal tastes and inclinations. In so doing, I fear that mass personalisation may foster the rise of widespread narcissism. Unlike consumerism, within a culture of narcissism, it is not the acquisition of material possessions per se that is valorised, but rather that of satisfying the individual's idiosyncratic tastes, preferences, views and inclinations. Because data-driven profiling enables mass personalisation well beyond the sphere of commercial consumption into the larger political and social sphere, I worry that it will foster and legitimise an attitude of narcissistic self-obsession, with detrimental effects not only for the individual but for society more generally.

In particular, I worry that the rise of mass personalisation may precipitate the rise of mass narcissism, and which might be so corrosive of social solidarity, and so loosen our social bonds to such an extent that the very nature of our collective character as a moral and political community might be called into question. Over time, the aggregate and cumulative effect of increasingly narcissistic tendencies and traits amongst individuals may fuel a larger narcissistic culture in which individuals come to believe in the validity of their own self-obsession and entitlement,

³⁸ Draper & Turow, *supra* n.8.

³⁹ Alain De Botton, *Status Anxiety* (Knopf Doubleday Publishing Group 2008).

fostering a narcissistic culture in which the individual's subjective needs, perceptions and desires are celebrated whilst concurrently fostering social isolation and undermining a sense of shared identity and collective purpose.⁴⁰ Viewed in normative terms, personalisation strategies focus on isolating and highlighting *differences* between individuals in ways that emphasise *particularity over generality*. This is the core logic underpinning market segmentation and data-driven personalisation strategies, which seeks to identify particular features about the targeted individual that are predictively identified as highly correlated with particular offers or messages, and thus more likely to be of direct relevance and particular interest to the individual in question. Yet, as Frederick Schauer has observed, practices based on generality and generalisation may have important community-building functions, noting that

'...at times, though certainly not at all times, generality and generalization may bring positive virtues. Generality, by excluding local variation (and not just in the geographic sense of 'local'), can serve important binding and community-creating functions... the rise of Europe as a power, a force, an idea, and a community has been a positive development, even if that positive development has come at the cost of some loss of local identity and national autonomy. By imposing uniformity in the face of diversity, harmonization in the face of difference, and consistency in the face of variance, generality flattens or dampens the variations among people, places, groups and events, often to the discomfort of those whose particular claims and situations are flattened, but often, as in Europe, to the benefit of the larger group.'⁴¹

It is this social binding function that inheres in, and is essential for, the idea of a community and which may provide one of the strongest arguments in favour of equality. As Schauer puts it, 'If we are not a community of equals, perhaps we are equals because we are a community.'⁴² He continues:

'The affinity between community and equality stems from our recognition of the real bite of the idea of equality. That bite emerges not from the fact of descriptive equality, but...from the fact that the most of the interesting exemplars of equality exist against the background of important inequalities. Rather than being important because it treats like cases alike, equality becomes important precisely because it treats unlike cases alike...

Even once we have recognised this fact, it remains puzzling for many people why we would want to treat unlike cases alike. Indeed, most of the sceptical writing about equality has maintained that equality for equality's sake has little to be said for it, and that typically what masquerades as equality is something else entirely. But perhaps there is something to be said for equality for equality's sake, and perhaps that something is closely related to the idea of community. That is, perhaps equality for equality's sake is important precisely because the level or Procrustean effect of equality for equality's sake is just what creates communities. To abstract away from our differences is to bring us closer together and to generate a focus on our similarities. When we ignore or abstract away from our differences, we necessarily increase our emphasis on shared standards and equal treatment...

When community is pursued, one important way of pursuing it, and one consequence of its pursuit, is the emergence of a generality whose effect, in turn, is to treat unlike cases alike in the service of community.'⁴³

⁴⁰ Robert D Putnam, *Bowling Alone* (Simon & Schuster 2001).

⁴¹ Frederick Schauer, *Profiles, Probabilities and Stereotypes* (Harvard University Press 2003) 289.

⁴² *Ibid*, 296.

⁴³ *Ibid*, 296-298.

By shifting the balance of values away from the collective interest, and its associated commitment to generality, in favour of particularity, mass personalisation might thereby weaken our commitment to the principle and practice of equality. In so doing, my fear is that it will weaken and dilute the communal and solidaristic character of our societies, such that our treatment of individuals persistently and systematically prioritises the differences between us rather than our similarities and shared features. In so doing, this necessarily entails a significant shift away from the community-building nature of generality, threatening to erode our sense of common purpose, community and solidarity while reinforcing social differences and exacerbating social stratification.

In expressing this fear, I am not suggesting that we should flatten and ignore all the differences between us, or that there is no individual and social value in seeking out and association with others whose tastes, behaviours, interests and political outlooks are similar to our own. As Turow has observed in relation to the increasing customisation of media environments that data-driven technologies have enabled, a healthy society needs both (a) 'segment-making media', which encourages 'small slices of society to talk among themselves'⁴⁴ and (b) 'society making media' - those which have the potential to promote conversation across segments⁴⁵. Each form of media has its own benefits and drawbacks for supporting a healthy society. While segment making media tends to offer their audiences a narrow set of views, society-making media has a track record of marginalising particular voices. Yet together both create the possibility for engagement within and across interest groups.⁴⁶ Similarly, healthy societies require both recognition of differences between individuals and groups which may warrant particular treatment, as well as recognition of their similarities and common features, in which equality of treatment is both necessary and justified. My fear is that by encouraging and legitimising an attitude of narcissism, in which individuals become so preoccupied with themselves and the immediate satisfaction of their particular interests, desires, needs and wants, there is a real risk that over time, the turn to mass personalisation might threaten our collective character as a community at all. A narcissist is, by definition, not merely preoccupied with him or herself, but considers himself or herself to be the ultimate object of admiration, without any interest or concern for others. If we, as a community of individuals, become increasingly narcissistic in our outlook, this might ultimately threaten the very nature of collective life as a community. Indeed, a community of narcissists is a contradiction in terms. It is a collection of isolated individuals, not a community of equals bound to each other through the bonds of a shared commitment to recognition of their common humanity, which demands both equality of treatment and an other-regarding attitude of mutual respect and self-restraint.

Conclusion

The breath-taking pace of technological change in recent decades is fuelling social transformations that are likely, in retrospect, to justify their description as 'revolutionary'. These include potentially profound changes to our modes and cultures of both production and consumption. In this paper, I have sought to reflect critically on the rise of mass predictive personalisation that data-driven service delivery in a hyperconnected age makes possible. While the tech and marketing industries champion their benefits by enabling users to enjoy a more 'meaningful' experience, I have suggested

⁴⁴ Joseph Turow, *The Daily You: How the New Advertising Industry Is Defining Your Identity and Your Worth* (Yale University Press 2012) 194.

⁴⁵ Ibid.

⁴⁶ Turow and Draper, *supra* n.8 at 1149

that all that glitters is not gold⁴⁷. In particular, I have identified five fears that the rise of mass predictive personalisation may portend for important collective values and commitments, which can be broadly understood as concerns about fairness and justice, on the one hand, and the erosion of social solidarity and loss of community on the other. In raising these concerns, I hope to provoke critical discussion and reflection that will motivate more penetrating research and place questions of this kind more firmly onto the policy, public and political agenda. It is vital that, in the face of so much hype, hope and fear associated our new data-driven technologies, we are actively involved in shaping our technological future, by asking ourselves - ultimately - what is the good digital life? Only then can we identify whether the vision of the good digital life currently offered to us Silicon Valley and its satellites is one that we willingly wish to embrace and, if not, to identify the social and technical governance mechanisms that are needed to reorient us, and to nurture and sustain the foundational values that are vital for individuals and communities to flourish in a data-driven age.

⁴⁷ Anthony Danna and Oscar Gandy, 'All That Glitters is Not Gold: Digging Beneath the Surface of Data Mining' (2002) 30 *Journal of Business Ethics* 373-386.