Data-Driven Personalisation and the Law – A Primer

Collective Interests engaged by Personalisation in Markets, Politics and Law

Uta Kohl, James Davey, Jacob Eisler*

Southampton Law School

Discussion Paper No: 1 (February 2019)

Call for Papers / Expressions of Interest

We will be holding a workshop on the topic of 'Data-Driven Personalisation in Markets, Politics and Law' on **Friday 28 June 2019** at Southampton Law School. This is an important emerging area of law that goes well beyond data protection law, raising questions for criminal law, consumer protection, competition and IP law, tort law, administrative law, human rights and anti-discrimination law, law and economics as well as legal and constitutional theory. To give the workshop focus and structure, this position paper provides a snap-shot of what we think about the topic or how we would frame it. We would like to hear your thoughts!

Should you be interested in disagreeing, elaborating, confirming, contradicting, dismissing or just reflecting on anything in the paper and present those ideas at the workshop, send us an abstract by **Friday 5 April 2019** (Ms Clare Brady C.L.Brady@soton.ac.uk). We aim to publish an edited popular law/social science book with the most compelling contributions after the workshop.

Introduction

When Netflix released *House of Cards* as its first original content in 2013, the series was accompanied by ten different cuts of its trailer. Subscribers were sorted depending on their prior viewing habits: did they like films with Kevin Spacey, or films directed by David Fincher or those with women as lead actors? As this technique has evolved, Netflix has attracted controversy for tailoring its advertising based on prior viewing profiles. One of its African American subscribers complained that *Like Father* was advertised to her with an image of two black actors with relatively minor roles in the comedy which by all accounts is a 'very white' film. Netflix denied that it engaged in racial targeting; it simply acted on what subscribers appeared to have liked in the past, stating: 'The artwork may highlight an actor that you recognize, capture an exciting moment like a car chase, or contain a dramatic scene that conveys the essence of a movie or TV show. If we present that perfect image on your homepage (and as they say: an image is worth a thousand words), then maybe, just maybe, you will give it a try.' (Deadline 2018)

In the lead up to the Brexit referendum in 2016, the official Vote Leave campaign spent £2.7m on targeted ads on Facebook. Vote Leave commissioned the Canadian company Aggregate AIQ to create a wide range of customised adverts, 1433 to be precise, that would combine a more or less explicit pro-Brexit message with one that tapped into the apparent passions, desires or fears of different Facebook users. A user profile that revealed an animal lover, or a tea enthusiast, or an environmental campaigner would be targeted with an ad that connected the Leave message with opposition to Spanish bullfights, an apparent EU threat to the 'cuppa' or EU-made obstacles in the protection of polar bears (BBC 2018).

Law enforcement is also increasingly relying on 'predictive policing' (Ferguson 2012; Elkin-Koren and Gal 2018). Predictive policing involves data sets from the criminal justice system on crime patterns overlaid with data sets on postcodes, demographics, traffic and weather patterns or sporting events to predict offences, and potentially also offenders and victims (New Scientist 2018). Predictive policing is justified on the basis of using scarce resources in the most cost-effective way, but has also been exposed for the biases within the prediction (Angwin 2016). West-Midlands Police along with eight other police forces has, since 2018, been running a National Data Analytics Solution project to predict violent offences based on a range of data sets of the criminal justice system: 'what we are going to do is look at a cohort of individuals who have already got convictions for gun and knife crime – many thousands of people – then we will let the data tell us what are the key predictive indicators on the journey they have been on to take them to the point where they stab someone or shoot someone.'(New Scientist 2018)

A comprehensive application of personalised law-and-order is emerging in China. In 2018, news media reported on the roll-out of China's social credit system that scores citizens to gauge their general compliance, legitimised on the basis that 'keeping trust is glorious and breaking trust is disgraceful' (China Copyright and Media 2015). Whilst at this point there appears to be a relatively uncoordinated collection of private and public social credit systems, the government scheme is expected to be fully operational nationwide and mandatory by 2020, and feed off the data of corporate credit systems, such as Alibaba's Sesame system. In parts, it already does so (Wired 2018). What is striking about the conception of the Chinese system is the extensiveness of the personal data collected and of the law-and-order purposes for which it is used, *and* the apparent lack of a relevant connection between the data collected and the purposes. Individuals receive a score based on data from social media activities, health records, insurances, private messages, financial data, gaming activities, smart home statistics, preferred newspapers or shopping histories, in addition to government data on serious crimes and anti-social misdemeanours such as bad driving, or smoking in non-smoking zones (Business Insider 2018). Low scoring individuals have been banned from flights and trains or had their internet speed throttled, they may be denied the best jobs or their children refused entry in good schools or universities. High scores come with rewards, such as priority access to public housing, travel visas and job promotions. The system's carrot-and-stick approach is bound to make it advantageous for many, and popular with large sections of the population. The Chinese system illustrates bluntly how personal data can be used as a powerful instrument of control, whilst the non-specificity underlying the system comments on the homogeneity of data itself and on violation of fairness caused through its arbitrary use. China has taken big data to another level in its governance agenda, but in liberal democracies, more specific personal scoring schemes based on large data sets are also increasingly driving decisions in civil life - from hiring to firing, to credit rating and insurance decisions – and not always without a backlash. Amazon abandoned its project to build a recruitment AI tool after it was shown to be biased against women (Business Insider 2018).

The above scenarios cover very different spheres of life - from light-hearted entertainment and service delivery within the video streaming market, to matters of intense political debate about the future of the UK within or outwith the EU and, finally, to law enforcement in the UK and law-and-order administration in China. What these scenarios have in common is the data-driven profiling of consumers or citizens to deliver a customised or personalised service, advert or legal response. Personal data is interpreted in conjunction with big data, using sophisticated algorithms, to create a picture of who someone *is* based on who they *were* - their past preferences, activities, networks and behaviours - in order to make a future-oriented prediction of what they might like (i.e. which film), what might persuade them (i.e. which ad), how they might act (i.e. commit a crime or succeed in a job) or what they might deserve (i.e. promotion or public housing). The above examples problematize personalisation, but in fact it has become ubiquitous online - it underlies every social media engagement, many apps and platforms from entertainment to food (Search Engine Land 2017). This trend is going to intensify with the rise of the internet of things, including smart medicine, smart cars and smart homes that will bring many other convenient, efficient and useful applications for personal data (Kelleher and Tierney 2018). Only the very occasional mishap will make the threat and manipulative potential of personalisation apparent.

Terminology offers some insights on its acceptability. In the market place where predictive profiling is welcomed by users and offered by businesses as added value, it is referred to as 'personalisation', 'customisation', 'optimisation' or 'smart' technology. A personalised recommendation, service delivery or health or fitness app is more useful and desirable than a non-personal one. Yet, even in the market place 'tracking' or 'targeted advertising' connote the more predatory side of the personalised business model. Profiling becomes opaque and thus goes underground, un-mentioned, when it is unwanted or perceived as unfair, such a personalised pricing or the personalised exclusion from certain products, e.g. credit or insurance (Zuiderveen Borgesius and Poort 2017). In the political sphere, it is also problematic that Facebook sells user vulnerabilities, their passions or fears, to political interest groups for their taking and manipulation. Still, personalised news delivery is prevalent and popular (Eskens et al 2017, Turow 2013). In the legal and criminal justice context, the terminology of 'predictive policing' or 'crime analysis' aligns data-driven profiling with ordinary policing practices and make it sound unexceptional. Overall the terminology of tracking, profiling, predicting, being smart, personalising, targeting is instructive in capturing the steps within, and the attributes of, personalisation as well as its beneficial and controversial aspects.

Data-driven personalisation within markets, politics and law is a mixed bag of innovative, useful, popular and benign opportunities and dystopian realities. The workshop seeks to unbundle the phenomenon in these different spheres from a legal perspective, with a focus on the collective goods or interests that personalisation may serve or undermine, and the commonalities and differences of those interests across the market, politics and law.

The Background: Technology, Behavioural Science and Economic Drivers

Personalisation of online content essentially filters information to make it 'fit' the particular user - based on 'recommender systems'. These systems are generally either content-based systems or collaborate filtering systems or, most commonly, a hybrid of the two. A content-based filtering system collects data on the interactions of the particular user with the site (e.g. Like or Dislike buttons, browsing and purchase history etc) and makes predictions based on these historic interactions. In other words, it looks at the actions of the particular user over time. Yet, it is collaborative filtering where big data really comes into its own, as it relies on a vast amount of data about *other* users to identify compelling similarities between different users. Amazon's success story is based on its relatively simply collaborative filtering: 'Customers who bought this, also liked….' The same type of system has been employed on a much bigger scale by social media sites, using large data sets on behaviour, demographics and geolocation. According to Facebook its 'average data set for CF [Collaborative Filtering] has 100 billion ratings [historical ratings of like-minded people], more than a billion users, and millions of items.' (Facebook 2015).

Personalisation has increasingly shifted from data reflecting the express selections, reviews, evaluations or ratings by users towards behavioural data. Typically once Netflix introduced streaming video, its 'recommender system became less about how you rate media and more about what you actually consume. It doesn't matter if you give critically acclaimed foreign dramas five stars if you spend all your waking hours bingeing *The Ranch*.' (Thrillist 2017) Even simple acts by users, such as clicking on Facebook Likes, have – through the use of machine learning tools - revealed correlations between certain Likes and personal attributes of users, including sexual orientation, ethnicity, religious and political views, personality traits, intelligence, happiness, use of addictive substance, parental separation, age, and gender (Kosinski et al 2013). Since then behavioural scientists at universities have been able to identify, again with the help of machine learning tools, markers of depression analysing 45000 photos from 166 Instagram users, 71 of whom had a known history of depression (Reece at al 2017); similar results have been achieved in relation to Twitter feeds (Reece at al 2017). These insights have promising applications in the diagnosis of mental illness – much like other big-data led behavioural science research.

However, in the commercial and political context, the uses of behavioural data are more contestable. To start with, for commercially driven personalisation, there is prima facie no need to directly identify personal attributes such as depression, sexual orientation or ethnicity, but they may be identified by proxy or as an unintended consequence of detecting common behaviour. For example, if the machine learning tool detects that certain subscribers watch comedies with certain actors, it predicts that these subscribers will enjoy other films with those actors even if those actors have minor roles. To the machine learning tool and in fact to the business, it is entirely by the by that those subscribers and actors are African Americans. The inference is amoral and non-judgmental.[1] This is not to say that it should not be criticised (see below). Still, Netflix was probably truthful when it denied intentional racial targeting. From a commercial perspective, the core attribute of each user lies in him or her becoming responsive to the 'product' i.e. the ad, the link, a movie, a service or the site or platform itself. Whatever works for the user in this respect, works for the provider.

The user's response to the personalised service provides the machine learning tool with automatic and constant feedback loops. Lower or higher engagement with the personalised service is fed back into the system and used to update what is presented to the user and like users. The machine learning tool is learning all by itself and this in fact means that even the initial designers of the tool cannot know the 'reasoning' behind the inferences or decisions made by it (Yeung 2018).

Personalisation of online content has been justified in terms of alleviating information overload by limiting all available information to what is 'relevant', 'interesting' or 'useful' to the particular user. The technology is pre-judging the information to make a selection that fits. The economic driver behind personalisation is profitability. For social media, personalisation increases the 'stickiness' of the platform; the longer a user spends on it, the more he or she will be exposed to its advertisement which is the main source of income for 'free' sites like social media, search engines, media sites etc. By implication, the success of companies is now measured by their capacity to capture human attention within the 'market for eyeballs' (Yeung 2018).

On a wider economic scale, the changes in production and consumption have been described as a shift from industrial capitalism to 'surveillance capitalism' (Zuboff 2015): mass production, or the production of identical units at scale, is replaced by mass personalisation, or the delivery of personalised services at scale (Yeung 2018, Deloitte Consumer Review 2015). This appears to be overstating the case considering that often personalisation only goes so far and the final product remains an identical product for mass consumption. Taking Netflix's *House of Cards* as an example, the trailer was 'personalised' to 13 stereotypes whilst the series was the same for everyone. Having said that, the Brexit ad campaign went further with 1433 stereotypes, and indeed for many industries like banking, insurance or utilities, big data is likely to make the transaction costs of highly personalised prices very low and thus feasible. This would allow for exceptionally precise price discrimination. Such practices may be justified on efficiency and user-pay grounds but also raise concerns regarding the desirability of powerful economic actors maximising their extraction from weaker, isolated actors.

The basic idea of identifying correlations in the activities and behaviours of users from large data sets, using machine learning tools, and comparing these with known users, also underlies predictive policing. Only here inferences are drawn from police databases and compared with the data of individuals that have already come into contact with the criminal justice system. For example, the project led by West Midlands Police draws on local and national police databases accounting for 5 million individuals, and has identified, through a machine learning system, 1400 markers relating to criminality, 30 of which are considered especially weighty. The objective here is that through automated matching of markers to the data of individuals known to be the police, it can identify those who are on a trajectory of violence and suggest interventions by social services (New Scientist 2018). Such carefully tailored enforcement maximises scarce public resources but, as in the domain of price personalisation, also allows a powerful actor (the state police apparatus) to exert enormous control over fragmented less powerful actors (individual citizens) and raises questions of due process and equal treatment under the law (see further below). Predictive policing also highlights that the use of artificial intelligence is entirely dependent on having relevant data about the target and that inequities may be created by having or not having a relevant digital identity.

The Law: Positioning AI-Driven Personalisation within Law and the Legal System

AI-driven personalisation appears to be, in the first instance, a matter only of data protection law: it relies on the collection of personal data which is processed to create a user profile in order to draw inferences and make predictions for a variety of purposes. Yet, there is more to it. For a start, as the primary focus of data protection law is to protect individual information privacy (balanced against the need for necessary and useful data collection and processing), it is relatively blind to other interests and rights that lie outside or at the margins of its natural or self-imposed remit. For example, it has been shown that freedom of speech, most importantly the right to receive information, is potentially strongly impacted – both positively and negatively - by personalised news media, and by implication the underlying values of democratic participation, truth, social cohesion and self-fulfilment it seeks to protect (Eskens et al 2017, Sunstein 2007, 2013). Thus freedom of expression ought to be debated in the personalisation context other than as a side note to data protection law.[2] Unlike Art 8 on the right to privacy, the jurisprudence on Art 10 of the European Convention of Human Rights has involved a wide focus on the public interests protected via the right to freedom of expression. This is not to say that privacy does not also protect important public goods (Solove 2007), but simply that individual privacy has not traditionally been articulated and justified in such public good terms nor in terms of possible privacy interests of groups (Taylor et al 2017).

Furthermore, even where data protection law offers some limited protection against purely automated decisions,[3] the newly introduced derogation in the form of consent is likely to make it even less effective than its already underused predecessor (Bygrave 2017). In any event, the review of automated inferences, predictions or decisions lies outside data protection law. The CJEU and the Advocate General 'have consistently restricted the remit of data protection law to assessing the legitimacy of the input data undergoing processing… [but not] the accuracy of decisions and decision-making processing involving personal data' (Wachter and Mittelstadt 2019). Essentially data protection law focuses on 'inputs' rather than 'outputs'. Review of the outputs belongs to other areas of the law. For example, the right to equality and non-discrimination in the employment context might provide relief. Equally, personalised advertising or pricing raise questions for advertising standards and may require a radical re-think of consumer protection laws. These sectoral laws, in turn, would, have to negotiate with or around trade secret law that protects the algorithms behind the personalised decisions, adverts, prices or products.

In so far as predictive profiling is used within the legal system itself, such as law enforcement in the criminal justice system, it also engages different procedural and substantive laws. Again, data protection law has little to offer towards deciding, for example, when predictive profiling would satisfy the 'reasonable grounds' condition of the police stop and search power (Ferguson 2012) or justify diverting the offender from custody to a rehabilitation programme (Wired 2018). This is something that needs to be negotiated within criminal law taking account of human rights, such as the right to a fair trial. Along similar lines, in civil law it has been asked whether the 'reasonable person' standard may be personalised to the particular defendant, in the same vein as negligence already qualifies reasonableness in light of some of his or her actual attributes e.g. the reasonable carpenter – although one wonders where the data sets would come from, short of adopting a Chinese-style digital identity model. In short, although data protection law is important for initially overseeing the collection and processing of personal data, it does not and cannot deal with the myriad interactions between data-driven personalisation and the legal system.

Still, what emerges from within the necessarily bounded perspective of data protection law is insightful. The GDPR delivers personalisation to companies on a golden plate. It does so, first, by shifting the prerequisite for more expansive uses and reuses of personal data from anonymisation to 'pseudonymisation'. The new concept of pseudonymisation is a half-way house between data that directly identifies the subject and anonymous data, by referring to data that identifies the subject, but only indirectly via additional information.[4] The reason behind the shift lies in the fact that, although anonymised data is effective in protecting privacy, much analytical value and utility of the data is lost through anonymisation (Ohm 2010) - as particularly relevant for personalisation purposes. So the GDPR makes some, but not all, of the processing privileges previously only enjoyed in relation to anonymised data to pseudonymised personal data.[5] More uses for less protected personal data.

Second, the GDPR facilitates personalisation by making the collection and processing/use of personal data essentially a matter of informational self-determination, i.e. notice-and-choice and the right to information.[6] (Wachter and Mittelstadt 2019) The GDPR's reliance on informational self-determination is instructive about the Regulation's deeper design and outlook in three ways - as embodied in three different fictions. First, consent and the implicit assumption of choice is fictitious, if not wilfully blind, in the current information landscape that is subject to powerful network effects and dominated by the Big Five (Google, Facebook, Amazon, Apple and Microsoft) who share data freely amongst their many tentacles. Whilst it is certainly not the function of data protection law to ensure a diverse and competitive market, designing a data protection regime that assumes such a market, when it is blatantly absent, suggests that consent was, from the outset, used as a convenient legitimising device. Second, it has not escaped behavioural economists that personalised algorithms feed off the fact that much 'individual decision-making occurs subconsciously, passively and unreflectively rather than through active deliberation' (European Commission 2018, Kahneman 2012) and that users can be nudged in subtle and insidious ways in line with a business strategy. By the same token, Terms & Conditions or Privacy Consent notices that rely on the rational actor fiction are ignored by most online users most of the time. This has been known for some time (Nordhausen Scholes 2012, Solove 2013, Ben-Shahar and Schneider 2014, Busch 2018). Whether forcing companies to provide simpler notices that make opting-out of personalisation realistic remains to be seen. [7] Third, an emphasis on information self-determination powerfully suggests to users - and the disinterested bystander - that all that is at stake when users agree to the collection and processing of their data is their own personal interest, and nothing else. There is no suggestion that the collection of data might matter to the provider over and above the particular user.[8] Yet, in fact it does so very much. The data enters the big pool of data that is the essential building block for profiling everyone else - as individuals and groups. There have been suggestions that the value of this personal data *to* companies could be monetised (Malgieri and Custers 2018), but such monetisation misses the point that important collective public goods at stake, such as deliberative democracy, ought not to be for sale. Having said that, this is in fact already an unintended consequence of users exchanging their data for free online services (Sunstein 2007). As a thought experiment, one might think about personal data as the new plastic where each user contributes to the pollution of the information landscape with their personal data – even if the source of the problem is structural. Thus here we have a classic tragedy of the commons, or collective action problem. Admittedly, this puts a negative spin on big data and its current usages. The more general point is that the regulatory emphasis in data protection law on informational self-determination makes it difficult to escape personalisation and also obscures the collective interests that are engaged when users are asked for choices on their personal data.

Finally, personalisation has also raised jurisprudential questions and matters of constitutional theory about whether a state-centred understanding of law and regulation as scripted in statutes and judgments backed by a monopoly of force remains adequate in the information age when conventional law is either trumped or simply rendered irrelevant by self-executing technological regulation or 'code' (Graber 2016, Yeung 2017, Nissenbaum 2011, Hildebrandt and Koops 2010). If 'Code has the effect of controlling conduct and is thus analogous to Hart's concept of primary rules… secondary rules would accordingly be rules that provide remedies for the defects of rules that are enshrined in the medium of code.' (Graber 2016) What shape these secondary, meta or constitutional rules might take remains to be seen, but techniques for evaluating algorithms, or their effect, are beginning to emerge (Demortain 2017).

The Problems: Common Themes about Collective Goods

The following suggests some of the general themes that emerge from the AI-driven personalisation across markets, politics and law. The themes are designed to reflect, and feed into, regulatory responses to personalisation in light of fundamental values and collective goods upon which the political, economic and legal regimes have rested. Some of these themes have already been articulated in different ways by other writers (Sunstein 2007, Graber 2016, Eskens et al 2018, Yeung 2018). The themes, as captured below, do not generally depend on exploitative practices to raise concerns, but would be exacerbated by such practices.

Personalisation & Stereotyping

The term 'personalisation' is misleading in so far as it suggests something like a tailored-made suit, i.e. one made to the unique measurements and wishes of the customer (Yeung 2018). AI-driven personalisation relies on the comparability or even sameness of the user or target with others in certain respects. Paradoxically, personalisation fundamentally denies individual uniqueness. Comparability is used to negotiate the diversity of humankind to identify groups to whom a product from a limited range of products can be matched e.g. a trailer or commercial or political advert or price, or a group that should attract police attention. Unlike 'normal' stereotyping, this form of stereotyping is not *a priori*, but follows a decision of relative comparability of tastes, preferences and behaviours. For that reason, it seems plausible that AI-driven personalisation could cut across established cultural, social or ethnic groups, and thereby be a liberating disruptive societal force. Yet, as the Netflix trailer advert shows existing cultural, social and ethnic groups frequently have a group digital footprint and personalisation unwittingly picks up on these footprints and reinforces them (Taylor et al 2017). This is also likely to express itself in the practice of personalised prices or personalised access to financial products where some consumer groups will find themselves paying above average prices due to lack of information or know-how or being altogether excluded from certain services (Yeung 2018). Furthermore, in so far as personalisation always uses data from the past and present as a blueprint for the future, it locks in existing inequalities and inequities between groups. As a matter of methodology AI-driven personalisation appears to be incapable of moving from the 'is' to the 'ought'. This explains not only why Amazon's recruitment tool based on past data dominated by successful men was biased against women, but possibly also why it was abandoned.

In the political sphere, the concern of stereotyping has been articulated through references to 'filter bubbles' and fragmentation of the political community into enclaves of like-minded people (Sunstein 2007). There is conflicting evidence and opinion on whether these 'filter bubbles' exist or not (Graber 2016) but, assuming they do, it seems a valid concern that they undermine deliberative democracy and social cohesion, as groups no longer talk to each other (Eskens et al 2017). In any event, an information environment that that caters for individual comfort levels undermines 'liberal democratic citizenship [that] requires a certain amount of *dis*comfort' and innovation more generally (Cohen 2013).

Even where personalisation or profiling is based only on past or present activities, behaviours or preferences of the particular user to draw inferences about his or her preferences in the future, stereotyping occurs and becomes a self-fulfilling prophecy. Rob Roy of 2018 will be the same or very similar to Rob Roy in 2019. There is no or little possibility for re-invention which would seem especially problematic in the criminal justice context.

Personalisation & Personal Autonomy and 'Choice'

Personal autonomy is engaged through personalisation in three distinct ways. First, as mentioned above, AI-driven personalisation locks personalities in time and perpetuates and reinforces them through the personalised approach. Second and related, the process through which personalisation occurs is largely beyond the control of the user as it is based on behavioural data drawn from subliminal or subconscious activity, rather than deliberate and expressed preferences. If Netflix disregards my express ratings in preference to my actual viewing practices, it prioritises my real self over my aspirational self – likely against my wishes. Third, most problematically, personal autonomy is engaged through personalisation because many users desire it and actively opt into the personalised narrowing of their future choices (regardless of the fact that companies make it very difficult to opt out of personalisation and the GDPR arguably facilitates it). So is it like sugar? Just because you like it, does not mean it is good for you. If the state makes interventions in an obesogenic environment, might it also make interventions to protect our collective interest in personal autonomy? Should the state override personal autonomy to safeguard personal autonomy?

Personalisation & Relevance, Reasoning and Communicating Reasoning

AT-driven personalisation identifies correlations (Mayer-Schoenberger and Cukier 2013) between behaviours of different users and draws inferences based on these; it does not offer an explanation underlying the correlation or for the inference. It does not explain, it simply shows patterns. To add to this, the complexity of the decision making by machine learning tools or algorithms that absorb and compute a huge amount of data points and improve their own 'understanding' through feedback loops means that their functioning is beyond 'systematic observation, identification of harmful effect, or investigation of their causes' (Wachter 2018) even though post-decision auditing may detect biases (Demortain 2017). Yet, decisions in law (e.g. a decision concerning custody) and, to some extent, in the market (e.g. personalised prices) require an explanation or justification as a matter of treating the subject with dignity, legitimizing the decision and holding the decision-maker to account. This more general intractable problem of algorithmic transparency and accountability has attracted significant academic attention (Andrews et al 2017, Yeung 2018) and rightly so, as it calls for Hart's secondary rules which are designed to repair the defects of primary rules (Graber 2016). General concerns of algorithmic decision-making are exacerbated in the context of personalised decisions, given that the absence of a common yardsticks hinders easy checks on fairness and equality.

The opaqueness of personalised decisions has raised a third area of unease about the reasoning behind them. Given the homogeneity of all data as 'bits' regardless of their underlying subject-matter, sensitivity or value, it is not just easy to merge and consolidate data sets (Wachter 2018), there is an almost irresistible temptation to do so, given that insights and inferences drawn from larger sets are invariably greater than the sum of its parts. Such consolidation and cross-usages, which occurs routinely within the private sector and between the private and public sector, brings into question the relevance of particular data to particular decisions (invoking reasoning along the lines of judicial review of administrative decisions) as well as the wider societal costs of allowing such consolidation and concentration of data (Elkin-Koren and Gal 2018, Solove 2007).

Personalisation & the Loss of a Common Frame of Reference

Data-driven personalisation removes the common frame of reference that allows groups and individuals to assess the relative fairness, legitimacy and integrity of a contract price, a sentence or a political view and collectively exert leverage over business or government. In commerce, it is the market place which reflects aggregate choices of consumers and providers, against which consumers can make buying decisions as 'rational economic actors' and these decisions act as a form of collective bargaining with the business. With personalised pricing and advertising, that common yardstick becomes elusive. In the political context, the joint space is, or was, the public sphere as occupied by traditional professional mass media which could be trusted (albeit imperfectly) to generate factually neutral accounts of political events that would stake out the parameter within which public opinions can be formed. In a world of personalised news, there are no common facts that bind arguments on either side together. Finally, in the legal sphere, it is the rule of law and equality before the law which provide the frame of reference against which the fairness, legitimacy and integrity of legal decisions can be judged. Personalised law, even if only at the level of law enforcement, is on a very fundamental level at odds with this common foundation as it promotes, at its very core, the idea that everyone is unequal and deserves an unequal treatment. Thus data-driven personalisation, whether intentionally or otherwise, works on the premise of 'divide and conquer' rather effectively (Yeung 2018).

Personalisation, Efficiency, and Technical Domination of Law and Society

Technical research on personalisation has primarily focused on increasing the efficiency and efficacy of algorithms (Adomavicius and Tuzhilin 2005, Ricci et al 2015). From this perspective, the ultimate goal in the implementation of personalisation systems is to ensure that consumers receive optimal recommendations from any personalised system they use. If achieved, personalised algorithmic systems could ultimately prove the 'best' mechanism for decision-making, it will displace less-optimised alternatives. In theory, personalised algorithms could produce a condition of 'bliss' for participants, as no alternative means of decision-making could improve personal welfare. Integrated into broader social systems, such techniques could yield optimal allocation of resources – ranging from entertainment to healthcare to romance – across society. Were such a system to be broadly achieved, it might even become self-perpetuating. From one futurist perspective, this is the ultimate culmination of market logic. Our society has already absorbed a dramatic shift in the volume of information routinely available to private parties, and a corresponding increase in the speed and inexpensiveness of processing. To an 18th century capitalist, we might appear to be so close to the state of bliss as to be indistinguishable from that final destination.

From a classical legal perspective, this trend poses little threat to either legality or social welfare. The dominant vision of welfare in law has emphasised efficiency through minimisation of transaction costs as the goal towards which legal systems should aspire (drawing from Coase 1960). This vision underpins much of free-market thinking; it also is highly compatible with a permissive attitude towards decision-making guided by personalisation algorithms, so long as participants consent to the use of such algorithms.

Yet the unflinching validation of personal decision-making by algorithmic function as optimal ignores the constructed character of the very system it produces. The idea that preference-satisfaction can be objectively optimised depends upon an external imposition of a universal set of values, a project that personalisation itself cannot provide. Personalisation cannot by itself validate itself as the central organising mechanism of society, but only reinforce itself as the dominant means of ascription of meaning. This is linked to a basic observation of systems theory: the fundamental property of systems is communication (Luhmann 1995). Actions as well as language that allow for exchange of information between actors are the universal feature of society, including law and economics. Communication is not merely a secondary property of social systems, but the very means by which values are instantiated. Algorithms advance and reinforce a particular set of values, and their elimination of transaction costs is simply the elimination of barriers to this reinforcement.

The triumph of choice guided by personalised algorithms would not be objective optimisation of welfare, but the valuation and self-reinforcement of a set of outcomes that can be achieved by a type of technical process. This is, in itself, unproblematic. Personalised algorithms, and the broader framework that has produced them, parts of the emergent social system (Hayek 1973). More problematic, however, is the apparent belief – particularly among those who are implementing such systems – that the efficiency maximisation achieved by personalised algorithms is a first-order good, and that the only necessary check upon them is the consent of participants. In short, the designers and refiners of personalised algorithms have generally adopted an attitude that input control alone is sufficient to alleviate any systemic concerns. Rather, the deeper challenge of increasingly effective personalised algorithms is that they will displace alternative forms of communication.

This concern has more bite when brought to the level of practical reality, and weaves together the themes of this proposal. As personalisation becomes increasingly effective and pervasive, it will grant enormous of power to a handful of designers and overseers of the technology; it will reshape the framework within which individuals make choices and understand their own freedom; it will displace social forms of communication with a technical form of communication; and it will populate individuals' own preference with tailor-made content. This will certainly produce remarkable and self-reinforcing social changes. Yet whether or not this comprises a moral good is a question that can only be asked by contextualising the outputs of personalisation.

Law comprises an alternative communicative system. It has the ability to impose barriers to certain types of conduct; to shift incentive calculations; to advance norms. It does so through its own logic of politically autonomous legislation and judge-made precedential reasoning. It thus comprises a mechanism that can both critique and discipline personalisation, both through direct interventions and through providing a basis for contextualisation, reflection, and critique.

Bibliography

Adomavicius, Gediminas and Alexander Tuzhilin, 'Personalization Technologies: A Process-Oriented Perspective', 48 *Communications of the ACM* 83 (Oct 2005).

Andrews, Benbouzid, Brice, Bygrave, Demortain, Griffiths, Lodge, Mennicken, Yeung, *Algorithmic Regulation* (Discussion Paper No 85, LSE, 2017)

Julia Angwin, Jeff Larson, Surya Mattu, Lauren Kirchner, 'Machine Bias – There's software used across the country to predict future criminals. And it's biased against blacks' (23 May 2016) *ProPublica* https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing

BBC, 'Vote Leave's targeted Brexit ads released by Facebook' (26 July 2018) https://www.bbc.co.uk/news/uk-politics-44966969

Omri Ben-Shahar, Carl E Schneider, *More Than You Wanted to Know: The Failure of Mandated Disclosure* (Princeton 2014)

Roger Brownsword, 'Law, Liberty and Technology' in Roger Brownsword, Eloise Scotford, Karen Yeung, *The Oxford Handbook of Law, Regulation and Technology* (OUP 2017) 41.

Christoph Busch, 'Implementing Personalised Law: Personalised Disclosure in Consumer Law and Privacy Law' (2018) https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3181913

Christoph Busch, Alberto De Franceschi, 'Granular Legal Norms: Big Data and the Personalization of Private Law (forthcoming in Mak, Tjin, Berlee (eds), *Research Handbook on Data Science and Law* (EE, 2018)

Business Insider, 'Why it's totally unsurprising that Amazon's recruitment AI was biased against women' (13 October 2018) https://www.businessinsider.com/amazon-ai-biased-against-women-no-surprise-sandra-wachter-2018-10?r=US&IR=T

Business Insider, 'China has started ranking citizens with a creepy 'social credit' system — here's what you can do wrong, and the embarrassing, demeaning ways they can punish you' (29 October 2018) https://www.businessinsider.com/china-social-credit-system-punishments-and-rewards-explained-2018-4?r=US&IR=T

Lee A Bygrave, 'EU data protection law falls short as desirable model for algorithmic regulation' in Andrews, Benbouzid, Brice, Bygrave, Demortain, Griffiths, Lodge, Mennicken, Yeung, *Algorithmic Regulation* (Discussion Paper No 85, LSE, 2017)

China Copyright and Media, 'Planning Outline for the Construction of a Social Credit System (2014-2020)' (25 April 2015) https://chinacopyrightandmedia.wordpress.com/2014/06/14/planning-outline-for-the-construction-of-a-social-credit-system-2014-2020/

RH Coase, 'The Problem of Social Cost' (1960) 3 *Journal of Law and Economics* 1

Julie E Cohen, 'What Privacy is For' (2013) 126 *Harvard Law Review* 1904

Deadline, 'Netflix's Use of Artwork Personalization Attracts Online Criticism' (22 October 2018) https://deadline.com/2018/10/netflixs-artwork-personalization-attracts-online-criticism-1202487598/

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

David Demortain, Bilel Benbouzid, 'Evaluating predictive algorithms' in Andrews, Benbouzid, Brice, Bygrave, Demortain, Griffiths, Lodge, Mennicken, Yeung, *Algorithmic Regulation* (Discussion Paper No 85, LSE, 2017)

William H Dutton, Bianca Reisdorf, Elizabeth Dubois, Grant Blank, 'Search and Politics: The Uses and Impacts of Search in Britain, France, Germany, Italy, Poland, Spain, and the United States' (1 May 2017) *Quello Center Working Paper* No. 5-1-17, https://ssrn.com/abstract=2960697

Niva Elkin-Koren and Michal S Gal, 'The Chilling Effect of Governance-by-Data on Innovation' (2018) 86 *Chicago Law Review*

Sarah Eskens, Natali Helberger, Judith Moeller, 'Challenged by news personalisation: five perspectives on the right to receive information' (2017) 9 *Journal of Media Law* 259.

European Commission, *Behavioural Study on the Transparency of Online Platforms* (Final Report, 2018)

Facebook, 'Recommending items to more than a billion people' (2 June 2015) https://code.fb.com/core-data/recommending-items-to-more-than-a-billion-people/

Christoph B Graber, 'The Future of Online Personalisation: Technology, Law and Digital Freedoms' (2016) i-call working papers No 2016/01, University of Zurich

Andrew Guthrie Ferguson, 'Predictive Policing and Reasonable Suspicion' (2012) 62 *Emory Law Journal* 260

Hayek, F.H.A., *Law, Legislation, and Liberty, Volume I: Rules and Order* (University of Chicago Press, 1973)

Mireille Hildebrandt, Bert-Jaap Koops, 'The Challenges of Ambient Law and Legal Protection in the Profiling Era' (2010) 73 *Modern Law Review* 428

Daniel Kahneman, *Thinking, Fast and Slow* (Penguin, 2012)

John D Kelleher and Brendan Tierney, *Data Science* (MIT Press, 2018)

Michal Kosinski, David Stillwell, Thore Graepel, 'Private traits and attributes are predictable from digital records of human behavior' (2013) PNAS https://www.pnas.org/content/pnas/early/2013/03/06/1218772110.full.pdf

Luhman, Niklas, *Social Systems* (Stanford University Press, 1995) (John Bednarz, Jr. trans)

Gianclaudio Malgieri, Barrt Custers, 'Pricing privacy – the right to know the value of personal data' (2018) 34 *Computer Law & Security Review* 289

New Scientist, 'Exclusive: UK police wants to AI to stop violent crime before it happens' (26 November 2018) https://www.newscientist.com/article/2186512-exclusive-uk-police-wants-ai-to-stop-violent-crime-before-it-happens/

Annette Nordhausen Scholes, 'Behavioural Economics and the Autonomous Consumer' (2012) *Cambridge Yearbook of European Legal Studies* 297

Helen Nissenbaum, 'From Preemption to Circumvention: If Technology Regulates, Why do we Need Regulation (and Vice Versa)?' (2011) 26 *Berkeley Technology Law Journal* 1366

Viktor Mayer-Schoenberger, Kenneth Cukier, *Big Data* (Eamon Dolan Book, 2013)

Paul Ohm, 'Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization (2010) 57 *UCLA Law Review* 1701

Eli Pariser, *The Filter Bubble: What the Internet is Hiding from You* (Penguin, 2011)

Andrew G Reece, Andrew J Reagan, Katharina LM Lix, Peter Sheridan Dodds, Christopher M Danforth, Ellen J Langer, 'Forecasting the onset and course of mental illness with Twitter data' (2017) *Nature* http://www.uvm.edu/~cdanfort/research/2017-reece-scirep.pdf

Andrew G Reece, Christopher M Danforth, 'Instagram photos reveal predictive markers of depression' (2017) *EPJ Data Science* https://epjdatascience.springeropen.com/articles/10.1140/epjds/s13688-017-0110-z

Ricci, Rokach, Shapira, Kantor (eds), *Recommender Systems Handbook* (Springer, 2015)

Frederick Schauer, *Profiles, Probabilities, and Stereotypes* (Harvard University Press, 2006)

Search Engine Land, 'Google shortcuts use personalization and rich content to build next-gen local discovery tool' (4 May 2017) https://searchengineland.com/google-shortcuts-next-gen-local-discovery-tool-274411

Search Engine Land, 'Google admits it's using very limited personalization in search results' (17 September 2018) https://searchengineland.com/google-admits-its-using-very-limited-personalization-in-search-results-305469

Daniel J Solove, 'I've Got Nothing to Hide' and Other Misunderstandings of Privacy' (2007) 44 *San Diego Law Review* 745

Daniel J Solove, 'Privacy Self-Management and the Consent Dilemma' (2013) 126 *Harvard Law Review* 1880

Cass R Sunstein, *Republic.com 2.0* (Princeton University Press, 2007)

Cass R Sunstein, 'Impersonal Default Rules vs Active Choices vs Personalized Default Rules: a Triptych' (2013) https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2171343

Linnet Taylor, Luciano Floridi, Bart van der Sloot, *Group Privacy – New Challenges of Data Technologies* (Springer, 2017)

Thrillist, 'The Netflix Prize: How a $1 Million Contest Changed Binge-Watching Forever' (7 July 2017) https://www.thrillist.com/entertainment/nation/the-netflix-prize

Joseph Turow, *The Daily You: How the New Advertising Industry is Defining Your Identity and Your World* (Yale University Press, 2013)

Sandra Wachter, 'Normative Challenges of Identification in the Internet of Things: Privacy, profiling, discrimination, and the GDPR' (2018) *Computer Law & Security Review* 436

Sandra Wachter, Brent Mittelstadt, Luciano Floridi, 'Why a Right to Explanation of Automated Decision-Making does not Exist in the General Data Protection Regulation' (2017) *International Data Privacy Law* 76

Sandra Wachter, Brent Mittelstadt, 'A Right to Reasonable Inference: Re-Thinking Data Protection Law in the Age of Big Data and AI' (2019) *Columbian Business Law Review* https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3248829

Wired, 'UK police are using AI to inform custodial decisions – but it could be discriminating against the poor' (1 March 2018) https://www.wired.co.uk/article/police-ai-uk-durham-hart-checkpoint-algorithm-edit

Wired, 'The complicated truth about China's social credit system' (21 January 2019) https://www.wired.co.uk/article/china-social-credit-system-explained

Karen Yeung, 'Algorithmic regulation: A critical interrogation' (2018) 12 *Regulation & Governance* 505

Karen Yeung, 'Five Fears about Mass Predictive Personalisation in an Age of Surveillance Capitalism' (2018) https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3266800

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

Frederick Zuiderveen Borgesius, Joost Poort, 'Online Price Discrimination and EU Data Privacy Law' (2017) 40 *Journal of Consumer Policy* 347

* Southampton Law School. Thanks to Andrew Charlesworth, Alun Gibbs, David Gurnham, Diane Rowland and Christopher Harding for helpful comments, reflections and provocations.

[1] Having said that, more overt biases are bound to be created by labelling the correlations, which may be done to enable third party advertisers, as in the case of Facebook's Advertising Tool, which offers a huge number of categories of people e.g. 'Emotional security', 'Stop Illegal Immigration' or 'Household income: top 5% of ZIP codes (US)'.

[2] Art 85 GDPR.

[3] Art 22 GDPR (explicit consent generally required for automated decisions based on profiling having a legal or similarly significant effect on the data subject).

[4] Art 4(1) GDPR.

[5] See especially, in relation to further processing beyond the original purpose: Art 5(1)(b) and Art 6(4)(e); further processing for scientific, historical and statistical purposes: Art 5(1)(b) and Art 89(1) GDPR.

[6] Art 5(1)(b) GDPR (data must be 'collected for specified, *explicit* and legitimate purposes'), Art 6(1) (data is lawfully processed where the 'data subject has given consent to.. [it] for one or more specific purposes', OR where 'necessary for the performance of the contract' or where it is in its 'legitimate interest'); Art 13 (data subject must be informed about the collection of personal data); Art 21(2) (the right to object to profiling for direct marketing purposes); Art 22 (explicit consent generally required for automated decisions based on profiling having a legal or similarly significant effect on the data subject). Note, the controller's interests may at times trump those of the subject.

[7] The Guardian, 'Google fined record £44m by French data protection watchdog' (21 January 2019) https://www.theguardian.com/technology/2019/jan/21/google-fined-record-44m-by-french-data-protection-watchdog

[8] There are some encouraging signs that this is being recognised by regulators, see e.g. The Guardian, 'Italian regulator fines Facebook £8.9m for misleading users' (7 December 2018) https://www.theguardian.com/technology/2018/dec/07/italian-regulator-fines-facebook-89m-for-misleading-users