

# Is it getting too personal? On personalized advertising and autonomy

Sebastian Jon Holmen

Roskilde University, Department of Communication and Arts, Denmark  
Sjonh@ruc.dk

Early View publication date: 16 December 2023  
DOI: <http://dx.doi.org/10.5324/eip.v17i2.5056>



This is an open access article distributed under the terms of the Creative Commons Attribution 4.0 International License, which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

*It has recently been suggested that personalized advertising is often more of an affront to a person's autonomy and thus more morally worrisome than its generic counterpart precisely because it involves or takes advantage of such personalization. This paper argues that central reasons put forward to support this claim are unpersuasive and that generic and personalized advertising should therefore be treated as morally on par in terms of their potential to undermine consumer autonomy. The paper then suggests that, if this is true, scholars who defend the existence of moral asymmetry between personalized and generic advertising in terms of their effect on consumer autonomy need to choose from among three argumentative avenues. However, none of these avenues is likely to be particularly attractive for a defender of the asymmetry.*

**Keywords:** Personalized advertising, Generic Advertising, Autonomy

## Introduction

### Generic

A retailer dealing in many different household appliances also sells one type of bread maker. However, sales of this type of bread maker are low, so the retailer hires a marketing consultant to aid in upping sales. The marketing consultant suggests a strategy that involves sending ads for the bread maker to random email addresses and, without making it more salient on the website than other products, offering the bread maker for purchase on the retailer's website. The retailer pursues this strategy.

### Personalized

A retailer dealing in many different household appliances also sells one type of bread maker. However, sales of this type of bread maker are low, so the retailer hires a marketing consultant to aid in upping sales. The marketing consultant suggests a strategy that involves sending ads only to people who – based on data such as their recent online behaviours and demographics – are believed to be the most

likely to buy a bread maker. In addition, the bread maker is also made more salient on the website than other products for these particular consumers. The retailer pursues this strategy.

Philosophers, social scientists, law scholars and others have recently argued that cases like *Personalized*, where efforts to sell a product are based on data points about specific consumers' or groups of consumers' preferences, habits and cognitive vulnerabilities may pose a threat to their autonomy (for overviews, see Tsamados et al. 2021; Mittelstadt et al. 2016; Taddeo and Floridi 2018). Indeed, some even argue that such *personalized advertising* is often *more* of an affront to a person's autonomy and thus more morally worrisome than its generic counterpart – precisely because it involves or takes advantage of such personalization. *Generic advertising*, on the other hand, is based on statistical information regarding *general* preferences, habits and cognitive vulnerabilities. At a general level, a person is autonomous if and only if (i) she has the ability to identify what she has good reason to do, (ii) is able to be moved by her own reasoning, and (iii) can act in accordance with these reasons (Buss and Westlund 2018). While many theorists disagree about how to more precisely specify the content of these conditions, conditions (i) and (ii) are related to what has been termed the *decisional* dimension of autonomy, that is, to the autonomous formulation and determination of what ends to pursue. Condition (iii) relates to the *practical* dimension of autonomy. More specifically, it stipulates that persons must to some extent be able to pursue the ends they have set for themselves.<sup>1</sup> It has been argued that marketing efforts like advertising can potentially compromise especially the decisional dimension of autonomy (see e.g. Anker 2020).<sup>2</sup> As we shall see in what follows, the arguments positing an ethically relevant difference between generic and personalized advertising in terms of their effect on autonomy, can be understood to involve conditions belonging to both the decisional and practical dimensions of personal autonomy. What I shall argue in this paper, however, is that central reasons put forward to support this claim are unpersuasive and that generic and personalized advertising should therefore be treated as morally on par in terms of their potential to undermine consumer autonomy. I will then suggest that if this is true, scholars who defend the existence of a moral asymmetry between personalized and generic advertising in terms of their effect on consumer autonomy need to choose from among three argumentative avenues. However, none of these avenues seems particularly attractive for friends of the asymmetry.

I do not mean, of course, to deny the general claim that – given a vast amount of data and the ability to process it – corporate actors are, for better or (perhaps more often) for worse, increasingly capable of predicting and influencing individuals' online and real-world actions and opportunities in numerous ways (Citron and Pasquale 2014; Pasquale 2015). Companies also surely utilize the same or similar data and algorithms in order to send us the product ads which we are most likely to respond positively to at times when they are predicted to have maximum effect. That is, when the ads are most likely to affect our behaviour by having us engage with and (perhaps) purchase whatever is on offer. Companies surely also use the data points to instantly redesign their website so that the products consumers are predicted to want are placed “in the front of the virtual store” (Pariser 2011: 8). Indeed, the predictions that the data make possible can also

be used to design influences that attempt to “lure you into real places” (Zuboff 2019: 216), such as when you pass them on the street. I argue that the presence of these and other technical features of personalized advertising are currently not able to ground the view that personalized advertising embodies something especially morally sinister due to its effect on consumer autonomy. To support this conclusion, the paper proceeds as follows. In Section 2, I argue that, contrary to what some scholars have argued, no moral constraint can reasonably be said to exist on predicting individuals’ behaviour without their consent. In Section 3, I argue against the view that that personalization’s more limited information flow constitutes a threat to consumer autonomy. Section 4 discusses the view that personalized ads are particularly likely to take advantage of our cognitive vulnerabilities. Section 5 draws some implications from the analysis.

A few comments before we begin. First, to personalize ads to individual consumers or a group of consumers, a vast amount of data about them must obviously be fed to whatever algorithm the company in question employs for these tasks. This has rightly led scholars to point to important privacy and other ethical issues related to the acquisition of the data on which targeted or personalized ads are based (Véliz 2020; Zuboff 2019). However, I shall not engage with these concerns here. Second, and relatedly, personalized ads may be morally problematic or even wrong for companies to employ for reasons not related to their effect on consumer autonomy. For example, perhaps personalized ads are wrongfully exploitative, or perhaps they (at least sometimes) cause consumers significant harm. Third, as stated, I am interested in whether personalized advertising is *more* problematic than more generic forms of advertising in terms of its effect on persons’ autonomy. Hence, issues that seem to be present in both types of advertising in equal measure will not be engaged with here. For example, I shall assume that both types of advertising are equally likely to have deceptive content or be motivated by morally dubious intentions on the part of the advertiser.

### **Predicting behaviour without consent**

The first autonomy-based concern specifically raised in relation to personalized advertising that we shall consider is based on the fact that a company that sends out such ads bases its product choices on the user’s past behaviour and a prediction of their future preferences. This was exactly how we imagined that the retailers would attempt to market their bread makers in *Personalized*. Why might one find this practice to be an affront to a person’s autonomy? One reason might be that individuals have not provided their informed consent to predict their future preferences. Such predictions would show improper respect for the decisional dimension of an individual’s autonomy. In their widely cited paper laying out the potential ethical issues involved in the collection and use of big data, Herschel and Miori, for example, claim that:

[...] with Big Data individuals are frequently represented as data points that are then used to manipulate what the person will view in the future. That is, information is presented to individuals online that Big Data calculations determine best reflects their projected preferences based upon their previous search and online page view history. This algorithmic manipulation *presumes the will of the*

*individual without their explicit consent.* (Herschel and Miori 2017: 34, my italics).

As I understand these authors, the crux of the matter is that individuals should be consulted before predictions about their future preferences are made, to show proper respect for their decisional autonomy. Many might find this view intuitively appealing because it tracks the views expounded in other contexts, such as in biomedical ethics (Beauchamp and Childress 2009). Showing the proper respect for individuals' decisional autonomy, according to this view, requires obtaining their informed consent about matters that affect them. Upon further scrutiny, however, it becomes clear that this view must be rejected, because a claim to the effect that there is a consent-based moral constraint on making predictions about persons' future preferences is false. By way of example, suppose that I go to a store to buy a pair of shoes. I never buy my shoes anywhere else, and I always buy shoes with laces. Suppose further that the shoes I choose this day do not come with shoelaces and these must be purchased separately. Lastly, suppose that, when I go to the counter to pay for the shoes, the employee at the counter – having observed my earlier laced shoe purchases on several occasions – presents me with a selection of shoelaces in the belief that I will prefer to wear shoes that can be laced. In Herschel and Miori's account, this would seem to be a morally dubious action by the employee because he has predicted my preferences (or, "presumed my will") without my consent. However, surely that cannot be right, and hence, the consent-based constraint should be rejected.<sup>3</sup> It should be rejected, because a moral constraint that invites us to consider acts that are clearly morally innocuous as morally dubious or bad – even wrong – is itself highly dubious. Some might claim that this conclusion is premature because consent-based constraint applies only to algorithmic predictions, not the kind of prediction my case relies on. However, this seems *ad hoc*, absent a cogent explanation of what makes these two types of predictions relevantly different, and I cannot think of what would comprise such an explanation.<sup>4</sup>

In sum, if what I have argued in this section is true, then personalized ads are not more morally dubious qua more autonomy violating than generic ads because they assume the will of the person to whom the ad is served. This is so because the view that there is a consent-based constraint on making predictions about a person's future preferences is false.

### **Personalization limits information and options**

The second objection to personalized ads also turns on the personalized selection and presentation of information that it involves; however, it highlights that ads tailored to a person or group will result in diminished personal autonomy because of the information such ads *do not* offer regarding other purchasing alternatives.<sup>5</sup> Obviously, all ads, whether generic or personalized, likely leave out some information about product alternatives, but the narrowing of alternatives might be more pronounced when the range is customized based on data about a consumer's preferences or search history, for example. To provide a somewhat crude example, if I buy myself a new pair of golf clubs online, I might in the future be presented with ads for other types of golf equipment rather than, say, ads attempting to sell me rock-climbing gear. Alternatively, I might be presented only with ads for the particular brand of golf clubs that I have purchased rather than product

alternatives. This filtering of information has had some scholars sound the alarm. Jeannie Marie Paterson et al., for example, suggest that

[...] by removing alternative options from consumers' sight, targeted advertising narrows their opportunities for choice. This means consumers are making decisions from a position of less than full information, undermining the preconditions for the exercise of autonomy. This may well lead to a reduced number of and variation in the overall options presented to consumers. In other words, consumers may be constrained in their own echo-chambers of advertising that constrain their world view on the basis of their constructed digital profiles. (Paterson et al. 2021: 10)

In a similar vein, Eliza Mik (2016) has argued that when the range of product options that a consumer is presented with is personalized, “[...] his autonomy is limited as he is not given the opportunity to choose from – or become aware of – the full range of available options” (p. 21). I think we can discern two somewhat overlapping but ultimately distinct autonomy-related concerns from these quotes, and it is not clear to me whether Paterson et al. and Mik have both or only one of them in mind. On the one hand, these authors seem to take issue with the fact that personalizing ads is likely to worsen the available level of information about product options, information that could be important when consumers decide which product best aligns with their ends. Call this the *disclosure objection*. On the other hand, it seems that the mere fact that fewer action alternatives are presented to consumers motivates them to raise a moral red flag in the name of autonomy. Call this the *option objection*. Let us consider each in turn.

The *disclosure objection* derives some plausibility from the idea that disclosure of information relevant to a person's medical decision-making, for example, is often highlighted as a key to respecting their autonomy (Beauchamp and Childress 2009: 121f). It is obvious why this is the case; without information about different medical courses of action, it is impossible for a patient to judge which one aligns best with his or her preferences, values or long-term goals. Indeed, and perhaps more relevant for the present context, Thomas Anker (2020) has persuasively argued that failing to disclose “[...] information that is relevant, proportionate, sufficient and understandable to the average, targeted consumer [...]” (537) about a particular product fails to respect a consumer's decisional autonomy. Paterson et al. and Mik may hold that something similar is true in regard to the range of commercial products presented.<sup>6</sup> That is, personalized ads may often disclose too little information for a customer to make an informed choice about what product best aligns with their ends, such as their needs, preferences or values. However, I see at least two reasons why we should be sceptical of the *disclosure objection*. First, supposing that the personalized ads are well-informed, the products offered would likely be ones that do, in fact, align with the preferences and values of the receiver. Indeed, part of the very motivation for personalizing ads in the first place is presumably to attempt to offer us exactly the product(s) that are a good match for us, on the presumption that we will then more likely buy said product(s). Second, insisting, as Paterson et al. do, that “[...] making decisions from a position of less than *full information* [...]” (Paterson et al. 2021: 10, my italics) undermines the exercise of personal autonomy goes too far. If full information was necessary for

such decisions, then the fact that a person lacked knowledge about product alternative B at the time of purchasing product A would mean that their decision to buy A was non-autonomous, and this would be so *even if they were aware of option alternatives C through Z*. The problem is that this sets the bar far too high for what autonomy requires, and it implies that most (all?) of our decisions fail to be autonomous since most (all?) of us lack knowledge of at least some possible option alternatives. However, if this is true, Paterson et al. and Mik have yet to offer us a plausible explanation for when personalized ads undermine consumer autonomy by restricting their knowledge of different purchasing options. One way that they might attempt to provide such an explanation, would be to build on the work of Thomas Anker (2020), which similarly focuses on information about particular products. That is, perhaps they could attempt to argue that information concerning the range of available products presented to a consumer via a piece of advertising need not be complete (indeed, a complete list of products may undermine autonomy by causing informational overload), but nevertheless must meet certain conditions in order not to undermine autonomy. For example, the range of products presented must be relevant and proportional, that is, the information can be processed by the average targeted consumer “[...] within a reasonable period of time” (Anker 2020: 533). While I believe there may be something to this way of developing the objection in regard to the presentation of information about product alternatives in general, it is not entirely clear to me whether this approach would succeed in driving a wedge between personalized advertising and generic advertising in terms of their putative effect on consumer autonomy. For such an objection to stick it would need to be argued that only personalized adds fail to meet the specified condition(s) (whatever they may more precisely be), and I for one cannot think of any plausible conditions that would provide the grounds from which to argue this.

Let us now turn to the *option objection*. This objection, as will be recalled, concerns the relationship between reducing the range of options available to a person and a reduction of their autonomy. This objection shifts the point of concern from the informational background that consumers make purchasing decisions against – that is, the conditions for the proper formation of autonomous choice regarding what ends to pursue – to one concerning agents’ ability to act to attempt to achieve an end they have autonomously set for themselves. That is, the focus is moved from the decisional dimension of autonomy to its practical dimension. According to a view popular among some autonomy theorists, an agent is practically autonomous only if he has “[...] adequate options available for him to choose from” (Raz 1988: 373; see also Hurka 1987). So, perhaps the concern that Paterson et al. and Mik have in mind can be formulated in the following way: (1) presenting an agent with a personalized and thus more limited range of product options is a way of (2) reducing his or her action alternatives and (3) reducing action alternatives undermines practical autonomy. However, although the *option objection* thus formulated does have some intuitive appeal, it should, in my view, be rejected. The reason is that the move from (1) to (2) is illegitimate.<sup>7</sup> Specifically, it is false that the action alternatives in question – that is, acquiring goods other than the ones offered by the personalized advertising – are reduced by the fact that a limited range of products is being presented to the agent. That is, in principle, the options to purchase alternative products remain completely unaffected by

personalized ads, viz. the fact that my being offered a type of shampoo by a personalized ad has no effect on whether I have the option to go out and buy another type of shampoo. Hence, agents' practical autonomy is not limited by personalized ads, and the *option objection* fails. However, perhaps this is too quick. Some might want to follow Hurka (1987) in arguing that to be practically autonomous, "[i]t is not sufficient for autonomous action that a person has many options open. He must, most obviously, know about the options" (367). If this view is correct, it implies that consumers' beliefs about their option alternatives are central to them being practically autonomous. On this basis, some commentators might want to argue that personalized ads could reasonably cause consumers to believe that they have no action alternatives (even if they do have these alternatives). Thus, even if what I argued above is correct, consumers' practical autonomy might still be undermined by personalized ads because consumers are then effectively barred from acting in pursuit of their self-chosen ends. I am sympathetic to the general view that agents' beliefs about their option alternatives matter for an assessment of whether they are practically autonomous. However, it is ultimately an empirical question whether personalized ads ever have the suggested radical effects on individuals' beliefs (i.e., believing that there are no available alternatives), a question for which, to the best of my knowledge, we do not yet have an answer. Furthermore, and more importantly, if personalized ads can affect practical autonomy by affecting consumers' beliefs about action alternatives, then personalized ads may sometimes enhance rather than impede practical autonomy. For example, if I receive a personalized ad offering me an alternative product Y to a product X that I already use, and I happen to believe that X is the only product available, then the offering of product Y enhances my practical autonomy by making me aware of an option alternative that was already there. If this is true, then on this interpretation of what practical autonomy requires, we would have another reason to reject the move from (1) to (2).

To summarize briefly, in this section I have argued that concerns based on the idea that consumer autonomy is undermined by the information that personalized ads leave out are unpersuasive. Specifically, two somewhat overlapping concerns that commentators have argued personalized ads may have – turning on the restricted background information for forming autonomous choice and the limiting of options for exercising autonomous choice – fail to establish that personalized ads undermine autonomy more than generic forms of advertising do.

### **Taking advantage of irrationality and vulnerability**

Another plausible reason for thinking that personalized ads may undermine autonomy is that they might be used to take advantage of the now well-documented tendency of humans to make biased and heuristic decisions (for an overview, see e.g., Kahneman 2011). Obviously, taking advantage of, say, framing effects, anchoring bias, or some similar well-known irrational mechanism to induce us to buy a product is a common and often critiqued (Crisp 1987; Sher 2011) generic advertising strategy employed by companies in both online and real-world settings. Hence, it may seem that there is nothing special about personalized ads in this regard. To be clear, this is not to say that taking advantage of biases to sell products does not challenge consumer autonomy (perhaps even to a morally dubious extent),

but rather that if such a concern is sound, it generalizes to generic types of advertising as well. Compare:

#### Generic\*

A retailer dealing in many different household appliances also sells one type of bread maker. However, sales of this type of bread maker are low, so the retailer hires a marketing consultant to aid in upping sales. The marketing consultant suggests a strategy in which the company develops a second deluxe and more expensive version of the bread maker, which the company always presents before it presents the regular bread maker in their online ads and on their website. The online ads are sent out to random email addresses, and other than the deluxe bread maker always appearing before the regular bread maker, the bread makers are not more salient on the website than other products.

#### Personalized\*

Like the case above, except that the online ads are personalized in that they are sent only to people who – based on data such as their recent online behaviours and demographics – are believed to be the most likely to buy a bread maker. In addition, besides the fact that the deluxe bread maker always appears before the regular bread maker, it is also made more salient on the website than other products for these particular individuals.

In both *Generic\** and *Personalized\**, the retailer is instructed by the marketing consultant to attempt to take advantage of so-called anchoring bias to boost bread maker sales.<sup>8</sup> If we believe that taking advantage of such bias in an effort to increase sales of a product is morally dubious qua autonomy undermining (or not), we would suspect that they are equally so. Is there any reason to reject this view? Some have claimed that there is.

Imagine, following Ryan Calo (2014: 1033), that the facts the retailer had about their customers included knowledge about *what* biases they are each most susceptible to and *when* they are most susceptible to making irrational decisions (say, because they experience decision fatigue around this time of day). For the sake of brevity, call this the consumers' *irrationality profile*. According to several authors, if the retailer succeeds in having a consumer buy the bread maker by leveraging his or her *irrationality profile*, the consumer's decision to buy it is not fully autonomous (Calo 2014; Paterson et al. 2021; Strycharz and Duivenvoorde 2021; Susser et al. 2019; Yeung 2017). I tend to think that this is correct. However, I am not convinced that the fact that an advertiser uses specific consumers' *irrationality profiles* to personalize ads makes such ads more morally objectionable in terms of their effect on autonomy than generic ads, which also take advantage of the biases we humans harbour. The generic marketing strategy pursued in *Generic\** also took advantage of a cognitive bias, as we saw above. However, several of the cited authors would likely disagree for one or both of the following two reasons. In my view, they are both unpersuasive.



## Transparency

What level of transparency is morally required for it to be morally permissible to attempt to take advantage of individuals' cognitive and affective biases has recently been a central point of discussion in relation to the ethics of nudging (Sunstein 2016; Thaler and Sunstein 2021). Furthermore, transparency matters for autonomy because, as Susser, Roessler, and Nissenbaum observe, if “manipulees are unaware that features of their choice environments have been intentionally designed to influence them, their capacity to (competently) deliberate is undermined, yielding decisions they cannot endorse (authentically) as their own” (Susser et al. 2019, p. 38). In line with this, Mik (2016) argues that “[a] greater threat to autonomy is posed by those technologies that influence consumer choices in an imperceptible manner” and further that “[t]he influence of advertising is [...] generally limited by its easy detection” (p. 16). Hence, the first reason that some people hold personalized advertising to usually be more of an affront to individuals' autonomy than generic advertising is that the harnessing of bias in the former is less transparent than in the latter. We can distinguish two ways to argue that transparency is lower in *Personalized\** than in *Generic\**.

First, we might worry that the *token interference transparency* (Bovens 2008) is lower in *Personalized\** than in *Generic\**. That is, it will be less clear to consumers that personalized ads may contain attempts to use their bias(es) to induce them to purchase a product, than is the case with generic ads. In the debate on the ethics of nudging, token interference transparency is usually recognized to exist if an announcement about the use of nudging for a particular purpose is made (Thaler and Sunstein 2021). However, there seems to be no reason to think *Personalized\** could not be accompanied by such a message if *Generic\** could.

However, and secondly, it may be suggested that personalized ads are generally less transparent than generic ads in another way. This is especially so because the specific bias(es) that personalized ads attempt to utilize are harder for consumers to detect or observe than those that generic ads employ. That is, the *type interference transparency* (Bovens 2008) of the biases harnessed in personalized ads will be lower. However, there seems to be no reason to think that this is generally true. Even if a company possesses extensive knowledge about a consumer's irrationality profile, and thus what biases this person is most susceptible to, the biases they attempt to exploit to make the consumer buy their product will presumably be no different in kind than those to whom generic advertisers cater. Hence, if consumers can (not) recognize the appeal to a particular bias or set of biases in personalized ads, then they should (not) be able to recognize such appeal in generic ads either, and vice versa. One way to challenge this claim would be to point out, following Karen Yeung, that personalized influences informed by big data (what she terms 'hypernudges') are, in fact, often less transparent than generic types of influences because “[...] the critical mechanisms of influence utilized by hypernudging are embedded into the design of complex, machine-learning algorithms, which are highly opaque [...]” (Yeung 2017: 124). To put it differently, personalized influences (including personalized ads) are less transparent in the relevant sense because the personalized algorithms serving us the influence are likely too complex for us to observe (or proprietary laws may forbid us from observing) what bias(es) it is attempting to exploit. Taken to its most extreme, the influence delivered by the ad-dispensing algorithm could by design not even *in principle* be type interference

transparent (Bovens 2008) by making it impossible for an observer to discern what the influence intention is. Yet, even if this is so, I believe we have reason to refute that the potential opaqueness of which biases personalized ads attempt to exploit is sufficient for showing that personalized ads are therefore a greater affront to consumer autonomy than generic ads. Specifically, it is not clear that opaqueness is a feature reserved for personalized ads. Suppose, for example, that an advertiser had in his or her possession a machine-learning algorithm that could automatically rewrite or rearrange a piece of an advertisement so that it catered to one (or a configuration of several) of the hundred cognitive biases we humans harbour before it was sent out to a random list of consumers. It seems plausible to say that the generic ads delivered by this algorithm would be just as type interference transparent as the personalized algorithms that Yeung has in mind. This would seem to show that the potential lack of type interference transparency is not a result of personalizing influences but of the technical makeup of the algorithms employed in dispensing influences, be it personalized or generic.

### *More autonomy violations*

The second reason why some commentators have claimed that personalized ads employing irrationality profiles when exploiting a bias are likely more autonomy undermining than generic ads attempting to exploit biases relates to the sheer number of consumers who will be exploited (Calo 2014). Why should we worry that this number might be higher for personalized ads than generic ads? Suppose that one of the data points the retailer uses in *Personalized\** is something like the irrationality profile described above and that the retailer has such a profile for each of the consumers it targets. Each consumer in *Personalized\** would then be influenced to buy the bread maker at the exact time and in the exact way that the profile suggests would be most effective. In *Generic\**, by contrast, all consumers would receive the same biased ad, and some might not be susceptible to this particular type of bias or be so to a lesser degree than others. The result is that more people in *Personalized\** would be ‘paired up’ with the type and timing of the bias they are most vulnerable to than in *Generic\**, and thus more people would likely have their cognitive pitfalls exploited (and thus their autonomy thwarted) in the former than in the latter case. As Calo (2014) puts it, if this becomes a technical reality, “the ethical and legal distinction between the ordinary and vulnerable consumer” would collapse (2014: 1033). This is a scary scenario indeed, and it gains plausibility from observations such as the one that Facebook’s algorithm is demonstrably able to detect when teens are in potentially vulnerable states, such as feeling ‘useless’, ‘anxious’ or ‘stressed.’ Facebook then shares this information with advertisers, although it is not clear whether this information is used to personalize ads to these teens; Facebook denies that it is (Machkovech 2017).

It is certainly plausible that if personalized ads lead to a greater number of autonomy violations than generic ads, this constitutes a good reason to consider the former morally worse than the latter on autonomy grounds. However, as Calo also recognizes himself, it is ultimately an empirical question whether companies in the future will be able to successfully construct and leverage individual persons’ irrationality profiles in personalized ads to an extent that makes them more effective than generic ads also attempting to take advantage of our cognitive shortcomings. However, it is worth noting that the effectiveness of personalized ads

in general and ads that leverage personal information, such as when and how each of us is most vulnerable, may face consumer backlash resulting in reduced effectivity that generic ads do not. Specifically, it has recently been found that if ads are recognized by individuals to be moulded on personal information, they are perceived as inappropriately appropriated or used in a setting they should not be; this reduces the effectiveness of the ads (Kim et al. 2019). It seems conceivable that a similar effect would be observed regarding data used to construct irrationality profiles, causing the effectiveness of personalized ads employing such profiles to decrease significantly. It is of course an open question whether consumers would themselves indeed be able to recognize whether data about their particular psychological shortcomings are being leveraged against them. However, given that regulators and theorists are increasingly calling for increased transparency concerning what data is used to personalize ads, consumers may be aided in this regard.

To summarize, this section has argued that reasons offered in support of the view that personalized ads attempting to take advantage of consumer biases—their irrationality profile—are likely to be a greater threat to consumer autonomy than generic ads attempting to leverage common biases are unpersuasive. In particular, it was argued that there seems to be no reason to think that personalized ads relying on person-specific psychological insights would be less transparent and thus more likely to thwart autonomy than generic ads. It was also noted that the idea of increased effectiveness of personalized ads employing irrationality profiles remains an open empirical question. Note, however, that my analysis does not support the conclusion that there is nothing morally dubious about advertising taking advantage of biases, but rather that any conclusions we make about cases such as *Personalized\** on autonomy grounds apply with equal force to cases such as *Generic\**. What is the implication of this conclusion? It seems to imply is that friends of the asymmetry have three options, but none of them seems particularly attractive.

First, friends of the asymmetry can hold that, regardless of what analysis may show, there is something especially intuitively problematic in terms of autonomy going on in cases like *Personalized\**. Our moral intuitions about a case may, according to some philosophers at least, give us *some* reason to believe it to be morally dubious, however, this is usually considered quite a weak reason in need of further argumentative backing. Hence, if this is the route the sceptic of personalized advertising chooses, their critique will arguably be without much force.

The second option is to hold *Personalized\** and *Generic\** and similar cases to all be morally dubious qua autonomy undermining. However, the challenge for this strategy is to demonstrate how these kinds of marketing strategies are relevantly different from other influences that similarly affect individuals' choices and behaviour (e.g., framing effects, interior design marketing, etc.), influences that usually do not give rise to autonomy objections.

The third option is to abandon the idea that something morally sinister in terms of autonomy is going on in cases such as *Generic\** and *Personalized\**. People, they may say, are constantly being influenced by others in their choices and behaviours without this prompting a judgement of decreased autonomy, and ads (personalized or generic) are just one among numerous innocuous influences. This route, however, is likely to flout people's moral intuitions – surely there is something

morally dubious going on in both *Generic\** and *Personalized\** – and would thus either need to show why these intuitions are mistaken or explain them by reference to some other morally salient fact about the cases.<sup>9</sup>

Which of these options are preferable? I do not currently have a good answer, but I suggest that future work consider this question to be the central one to confront.

## Implications

Where do these observations leave us? The paper set out to answer the question of what arguments may be given for considering personalized advertising to be more of an affront to consumer autonomy than generic marketing efforts. Three such reasons were identified, but all turned out to be unpersuasive upon further scrutiny. I may of course have overlooked some relevant autonomy-based consideration(s). However, if I have not, the analysis provided here suggests that there is little reason to believe that personalized advertising qua personalized currently poses a greater threat to consumers' autonomy than generic advertising. As highlighted above, however, this does not rule out the possibility that sound autonomy-based objections could be raised against personalized advertising, but rather that such objections should then be raised with equal force against generic types of advertising (and vice versa). This is important because it invites critics to argue for a general variant of one or more of the three options discussed above: (1) Deny with reference to intuition that personalized advertising and generic advertising are morally on par, thus offering a quite limited defence of their position. (2) Bite the bullet and argue that personalized advertising and generic advertising are equally morally dubious because they are equally autonomy undermining. This would, however, need to be accompanied by an explanation of why similar generic marketing interferences that seem intuitively morally innocuous should not also be considered autonomy undermining. (3) Abandon the view that personalized and generic advertisements – including instances that take advantage of bias – are an affront to autonomy. However, for this strategy to accommodate the intuition that there is something morally problematic about cases like *Personalized\** and *Generic\**, it would need to offer a non-autonomy-based rationale for why such cases are morally dubious. As also noted above, however, none of these options is likely to be attractive for someone wanting to oppose the use of personalized, but not generic, advertising on autonomy-based grounds.

## Acknowledgements

My thanks to Thomas Søbirk Petersen, Jesper Ryberg, Fatima Sabir, Michael R. Hyman and an anonymous reviewer for their valuable comments on an earlier version of the paper.

## Notes

<sup>1</sup> For the distinction between decisional autonomy and practical autonomy, see e.g. Pugh (2020).

<sup>2</sup> For a radically different approach to the concept of consumer autonomy based on mathematical theory as well as a comprehensive overview of previous studies, see Hyman, Kostyk and Trafimow (2022).

<sup>3</sup> It is worth noting that alternative interpretations of Herschel and Miori's objection – e.g. that it is the lack of the consumer's consent to the specific *use* of the prediction rather than the making of the prediction itself that raises moral concern or the that it is the *combination* of making and using the prediction without consent that should trigger our moral outrage – are vulnerable to similar objections. That is, they too prove too much.

<sup>4</sup> For example, the explanation cannot be that the employee's prediction is more likely to be wrong than data-driven predictions (or vice versa), since the quality of the prediction seems to be irrelevant to whether a prediction is assuming my will.

<sup>5</sup> This concern is similar to, and often explicitly built on, the one raised by Eli Pariser regarding the personalization of internet searches, social media content, etc. more generally (Pariser 2011). However, as I understand Pariser, he does not claim that internet users' autonomy is undermined by being in what he calls 'filter bubbles', but rather that being in such bubbles may be detrimental to individuals' well-being and democratic legitimacy. In my view, this is a much more plausible claim than the one we are about to consider.

<sup>6</sup> It is important to notice that in contrast to the view offered by Anker, the objection raised by Paterson et al. and Mik does not seem to claim that personalized advertising is more of an affront to autonomy because such advertising is more likely to disclose inadequate levels of information about *particular* products, but rather that the *limited range* of products presented to the consumer involved in such advertising is what grounds this concern. In my view, this is also the most plausible way to formulate the concern (although I shall shortly argue that it ultimately fails) because there seems to be no reason to believe there to be a relevant difference between generic and personalized advertising in terms of providing adequate information about the products they advertise for.

<sup>7</sup> I am also not convinced that the move from (2) to (3) is a legitimate one since it seems to me that the question of whether one is positively barred from pursuing an end is a binary question. If this is true, then reducing the range of options will not violate practical autonomy so long as at least one action alternative remains open, and such an alternative will always be open in regard to personalized ads (i.e., not buying the product). However, this is a contentious view that requires a more elaborate defense that can be given here, and thus I will simply accept the move from (2) to (3) as legitimate.

<sup>8</sup> The company William-Sonoma employed a somewhat similar strategy when attempting to raise sales of their bread makers. A similar pricing strategy is now also suggested to be employed by other companies such as Apple (Schwartz 2020).

<sup>9</sup> One such fact may be that personalized advertising is more likely to cause consumers harm than generic advertising inter alia because the former, but not the latter, will often generate a feeling of unease in consumers. I thank an anonymous reviewer for drawing my attention to this possibility.

## References

- Anker, T. (2020). Autonomy as license to operate: Establishing the internal and external conditions of informed choice in marketing. *Marketing Theory*, 20(4), 527–545. <https://doi.org/10.1177/1470593120926255>
- Beauchamp, T. L., & Childress, J. F. (2009). *Principles of Biomedical Ethics* (6th ed.).

- New York: Oxford University Press.
- Bovens, L. (2008). The ethics of nudge. In T. Grüne-Yanoff & S. O. Hansson (Eds.), *Preference Change: Approaches from Philosophy, Economics and Psychology*. Berlin and New York: Springer.
- Buss, S., & Westlund, A. (2018). Personal Autonomy. In E. N. Zalta (Ed.), *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/personal-autonomy/>
- Calo, R. (2014). Digital market manipulation. *George Washington Law Review*, 82(4), 995–1051. <https://doi.org/10.2139/ssrn.2309703>
- Citron, D. K., & Pasquale, F. (2014). The scored society: Due process for automated predictions. *Washington Law Review*, 89(1), 1–33.
- Crisp, R. (1987). Persuasive advertising, autonomy, and the creation of desire. *Journal of Business Ethics*, 6(5), 413–418. <https://doi.org/10.1007/BF00382898>
- Herschel, R., & Miori, V. M. (2017). Ethics & Big Data. *Technology in Society*, 49, 31–36.
- Hyman, M. R., Kostyk, A., & Trafimow, D. (2022). True Consumer Autonomy: A Formalization and Implications. *Journal of Business Ethics*. <https://doi.org/10.1007/s10551-022-05114-0>
- Kahneman, D. (2011). *Thinking, Fast and Slow*. Farrar, Straus & Giroux.
- Kim, T., Barasz, K., & John, L. K. (2019). Why am i seeing this ad? The effect of ad transparency on ad effectiveness. *Journal of Consumer Research*, 45(5), 906–932. <https://doi.org/10.1093/jcr/ucy039>
- Machkovech, S. (2017). Report: Facebook helped advertisers target teens who feel “worthless” [Updated]. *Arstechnica*.
- Mik, E. (2016). The erosion of autonomy in online consumer transactions. *Law, Innovation and Technology*, 8(1), 1–38. <https://doi.org/10.1080/17579961.2016.1161893>
- Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2), 205395171667967. <https://doi.org/10.1177/2053951716679679>
- Pariser, E. (2011). *The Filter Bubble - What the internet is hiding from you*. London: Penguin Books.
- Pasquale, F. (2015). *The Black Box Society - The Secret Algorithms That Control Money and Information*. Cambridge, Mass: Harvard University Press.
- Paterson, J. M., Chang, S., Cheong, M., Culnane, C., Dreyfus, S., & McKay, D. (2021). The Hidden Harms of Targeted Advertising by Algorithm and Interventions from the Consumer Protection Toolkit. *International Journal on Consumer Law and Practice*, 9, 1–24.
- Pugh, J. (2020). *Autonomy, Rationality, and Contemporary Bioethics*. Oxford, United Kingdom: Oxford University Press. <https://doi.org/10.1093/oso/9780198858584.001.0001>
- Raz, J. (1988). *The Morality of Freedom*. Oxford University Press. <https://doi.org/10.1093/0198248075.001.0001>
- Schwartz, E. (2020). Apple’s Goldilocks Strategy. *Econlife*. <https://econlife.com/2020/12/pricing-strategy/>. Accessed 19 April 2022
- Sher, S. (2011). A Framework for Assessing Immorally Manipulative Marketing Tactics. *Journal of Business Ethics*, 102(1), 97–118. <https://doi.org/10.1007/s10551-011-0802-4>

- Strycharz, J., & Duivenvoorde, B. (2021). The exploitation of vulnerability through personalised marketing communication: Are consumers protected? *Internet Policy Review*, 10(4), 0–27. <https://doi.org/10.14763/2021.4.1585>
- Sunstein, C. R. (2016). *The Ethics of Influence - Government in the Age of Behavioral Science*. Cambridge: Cambridge University Press. <https://doi.org/10.1017/CBO9781316493021>
- Susser, D., Roessler, B., & Nissenbaum, H. (2019). Online Manipulation: Hidden Influences in a Digital World. *Georgetown Law Tehcnology Review*, 4(1).
- Taddeo, M., & Floridi, L. (2018). How AI can be a force for good. *Science*, 361(6404), 751–752. <https://doi.org/10.1126/science.aat5991>
- Thaler, R. H., & Sunstein, C. R. (2021). *Nudge - The Final Edition*. New York, NY, USA: Penguin Books.
- Thomas Hurka. (1987). Why Value Autonomy? *Social Theory and Practice*, 13(3), 361–382.
- Tsamados, A., Aggarwal, N., Cowls, J., Morley, J., Roberts, H., Taddeo, M., & Floridi, L. (2021). The ethics of algorithms: key problems and solutions. *AI and Society*, (0123456789). <https://doi.org/10.1007/s00146-021-01154-8>
- Véliz, C. (2020). *Privacy Is Power*. London: Bantam Press.
- Yeung, K. (2017). “Hypernudge”: Big Data as a mode of regulation by design. *Information, communication and society*, 20(1), 118–136.
- Zuboff, S. (2019). *The Age of Surveillance Capitalism*. London: Profile Books, Ltd.