

The Moral Opportunities and Perils of Smart Wearables for Decisional Autonomy

By Niël Henk Conradie

Angewandte Ethik, RWTH Aachen

niel.conradie@humtec.rwth-aachen.de

Research project: InviDas

November 2021



Das Forschungsprojekt InviDas erforscht neue Formen der Datenvisualisierung und Gamification, um textuelle Datenschutzerklärungen verständlicher zu machen und Nutzer*innen von Wearables zu einer souveränen datenschutzrechtlichen Entscheidung zu ermächtigen. Der transdisziplinäre Forschungsverbund besteht aus Gesellschaft für Informatik e.V., Stiftung Digitale Chancen, Garmin Würzburg GmbH, RWTH Aachen, Universität Bremen und Otto-Friedrich-Universität Bamberg. Das Projekt wird vom Bundesministerium für Bildung und Forschung im Rahmen des Forschungsprogramms „Technik zum Menschen bringen“ gefördert. Mehr Informationen: <https://invidas.gi.de/>

Introduction

Smart wearable technologies - wearable technologies that incorporate the means for the collection and algorithmic processing of data in order to produce aim-guided and corrective output¹ - have seen an explosive growth over recent years, with current predictions being that this trend is set to accelerate (Seneviratne et al 2017; Dian et al 2020). This proliferation has extended across disparate domains, ranging from medicine and military applications to fitness and social technologies. As with any technology that promises to leave such a footprint on society, and particularly one that involves emergent technologies, these wearables² present opportunities for our benefit as well as moral challenges we cannot be negligent in addressing. Our purpose in this piece is to narrow down on one vital facet of our lives where the benefits and moral challenges raised by these technologies are especially salient: decision-making.

The structure of the discussion is as follows: first we will unpack the moral importance of decision-making, which involves understanding the relationship between decision-making and autonomy. In this first section we also introduce the important idea that different external influences on an agent can be autonomy-promoting or autonomy-reducing, and that there is a, at least pro tanto, moral obligation to support the former and combat the latter. In Section 2, we discuss the qualities and capacity of wearables that are likely to have the most significant and most novel impact on our decision-making abilities. The qualities we identify are: ubiquity, proximity, and convenience, and we show how they can both support and reduce autonomy. This impact on autonomy is usually the result of the most significant capacity possessed by these technologies: the capacity to facilitate cognitive offloading, a capacity that requires data collection. Finally, in Section 3, we provide the list of moral opportunities and concerns we have identified, along with suggestions for how we might take advantage of the former while avoiding the latter.

1. The value of decisional autonomy

In everyday life, we value our ability to make our own decisions. Though at times we may wish for others to take a difficult decision from us or throw up our hands in despair at our desultory options, it is exceedingly rare for us to warm up to the idea that someone or something else would make our decisions for us generally. Indeed, one longstanding tradition in philosophy holds that it is precisely the value of our making our own authentic choices - our self-governance or autonomy - from which most if not all moral value flows (Korsgaard 1996; Darwall 2006). Though there is no need to endorse this more extreme view, it is as uncontroversial as a claim about value is likely to be to say that autonomy is morally valuable. However, it is not necessarily the only moral value at stake when we consider decision-making. There are also consequentialist considerations regarding the outcomes resulting from a decision - the actual increase in well-being, for example, however this is defined - is also of moral significance, and so we have a moral reason to ensure that the “best” decision is reached, with best being that decision that maximises some well-being. Quite obviously

¹ What is meant here by “aim-guided and corrective output” will be discussed in more detail in Section 2.

² Though there is an important distinction between smart wearables and wearables more generally - the former being a subset of the latter - for convenience we will be using the term “wearables” as shorthand for smart wearables throughout this work. This is not to deny that non-smart wearables do not also present opportunities and challenges for decision-making, only that we will not (at least not directly) be addressing these here.

there can be instances of decision-making where ensuring the autonomy of a decision and maximising the well-being resulting from its outcomes can stand in tension, if not direct conflict. Without any intent of providing a final verdict for how to adjudicate these tensions, we do contend that within the limited space of *commercial* smart wearables, those designed, developed, purveyed, and supported by commercial entities such as corporations, a policy of *autonomy priority* should be followed. This is not a controversial contention, as it already follows from what is common in practice and law. Though states are sometimes empowered (rightly or wrongly) to violate autonomy in the name of well-being, we do not, nor should we, permit commercial entities the same permission. Similarly, we at times allow our friends or family to strongly impact our autonomy, though as we will discuss below this may not in fact be best understood as infringements on autonomy at all. If we adopt an autonomy priority position, then autonomy should be the first moral value we worry about, and downstream benefits of potential consequences should be considered only after autonomy impacts have been accounted for. In line with this, we will take *decisional autonomy* as the moral value of central importance from here on out.

However, we must be careful in identifying the relationship between making decisions and autonomy, since the two are not synonymous. To highlight this, we should consider that not all decisions are autonomous or *as* autonomous as others. If I am hypnotised to vote for a certain candidate in an election, for example, then this decision was not autonomous. If I am physically addicted to a narcotic and out of addictive compulsion choose to sell my most prized possession for a fix, this is (at least) *not fully* autonomous. To make a decision autonomously is for that decision to be:

1. The result of your own decision-making processes.
2. Guided by your authentic aims and values.
3. Without undue external influence.

Each of these elements requires expansion. For some decision, X, to be the result of an agent's decision-making processes, X must be the result of that emotico-cognitive mechanism or set of mechanisms that is responsive to *normative reasons*. Normative reasons, as the name implies, are reasons that we recognise as having normative force, they obligate us toward certain decisions or actions.³ These obligations can be *pro tanto*, and our reasons-responsive mechanisms are far from infallible, so this is not to say that everything we identify as a normative reason does in fact represent an obligation. This mechanism can thus be called a *reasons-responsive mechanism* and is how agents exert rational control over their behaviour. A key part of this responsiveness is also a responsiveness to veridical evidence, a fact ably identified by Niker et al. in their argument for evidence-responsiveness as a condition for autonomy (forthcoming). This requires that for an agent to exhibit autonomy it must be the case that her pro-attitudes must be open to correction in the face of veridical evidence.

This process of responding to reasons is, if autonomous, guided by our authentic aims and values. These aims and values will determine which states-of-affairs we take to be reasons to decide and act, since if a state-of-affair means that deciding or doing X will make it more likely that an agent can achieve some aim of hers, then she has reason to X. It might be thought, then, that autonomy is merely about ensuring that X is moving an agent closer to their aims. However, the qualifier "authentic" means that this is not *always* the case. The addict who craves a fix has "acquiring a fix" as an aim, and so has a *prima facie* reason to

³ At least in the practical sphere. There are also normative reasons in the theoretical sphere, these being reasons that obligate us to believe certain propositions.

decide to get one. But this was precisely an example where we did not want to say that the choice was wholly autonomous. The issue here is that not all of our pro-attitudes (our motivations in a general sense) are identified with our agency in the same way. To understand this most clearly it can help to consider cases of *akrasia*: situations where we judge that we *ought* to do X, but we do $\neg X$ instead. Anyone who has procrastinated or taken that extra helping of pudding will be familiar with this. Cases of *akrasia* are failures of rational self-governance, cases where our control over our decisions or actions is shown to be incomplete. The addict, if we presume that she would see her own addicted behaviour as all-things-considered undesirable, is being *akratic*. There is a rich debate as to how exactly the authenticity of our aims should be determined (Frankfurt 1998; Christman and Anderson 2005; Westlund 2009).⁴ We make no pretense that this work will resolve this dispute, and so we will not put forward an authoritative definition of authenticity, but will rather follow the advice of Justice Potter Stewart and follow the standard of “I know it when I see it”. However, where much might hinge on the exact understanding of authenticity at stake, we will provide a more detailed, if pluralist explanation of our positions.

The final requirement for a decision or action to be autonomous is that it must result without undue external influence. External in this case does not necessarily refer to outside the body, though often that will be the case. A tumour that is impairing an agent’s reasons-responsiveness would also count as an undue external influence, just as my friend bringing new information to my attention that informs my decision-making is an external influence, but not undue. The externality of an influence here means the degree to which the influence is outside the complex of my reasons-responsive decision-making mechanism and the aims and values that guide it. Our decisions are almost always (if not indeed always) being impacted by external influences. More troublesome to pin down than the meaning of “external”, is the precise content of “undue”. In many ways there is a similarity to trying to identify the content of authenticity in the previous condition. It is easy enough to identify influences that would count as undue: the afore-mentioned tumour, hypnosis, being carried away by a strong wind. Things do become more complicated where other agents are involved, but even here some influences seem patently undue. If an agent is subject to severe indoctrination and socialisation which brings about an alteration to her aims and values - without permitting her the reasonable opportunity to reflect upon these influences and modify her aims and values in line with the outcome of this reflection - and then she makes a decision in line with these indoctrinated values, this would not be a wholly autonomous decision. To unify these different sorts of influences under one banner, we propose the following provisional definition - the limits of which will be discussed in Section 3:

Undue external influence: an external influence that either modifies the reasons-responsive processes, the aims and values that guide these processes, or reduces the set of available options in such a way that this influence is not open to scrutiny and corrigibility by these same processes before the modification or its results.

With this rough account in hand, we can now consider what morally follows from it. If autonomy has moral value, then we have a *pro tanto* reason⁵ to promote it, and to avoid

⁴ For an excellent discussion of these intricacies, as well as an important original attempt to reconcile some of the existing positions, see the work of Niker et al. (forthcoming).

⁵ It is very important to keep in mind that the moral reason we have to promote and avoid reducing autonomy can potentially be outweighed by other moral reasons, and nothing we say here disputes that.

reducing or eliminating it. This applies both to our own autonomy and to that of others. Where the negative side of this duty is concerned, and where we are dealing with the autonomy of others, we get the straightforward obligation not to be or bring about an undue external influence on another agent. When dealing with our own autonomy it is an obligation to not bring about outcomes that damage our reasons-responsiveness mechanism or to reduce the authenticity of our aims or goals. When we consider the promotion of autonomy in ourselves, we have reason to improve the responsiveness of our decision-making mechanism, to strengthen the guidance exerted by our authentic aims, and to reduce or preferably eliminate undue external influences. Similarly, we have a reason to do the same for others. All these dimensions are important, but for the rest of the work we will be focusing on the promotion or reduction of autonomy in others.

To make matters more concrete: the promotion of autonomy in others can be achieved through a great variety of means, probably the most obvious of which is the provision of information relevant to a given decision. By equipping another agent with more or better information, one enables them to acquire knowledge about the facts relevant to a decision. This in turn makes it more likely that they will be able to recognise relevant reasons in the vicinity. Alternatively, highlighting options that better meet an agent's aims over others (assuming you have this knowledge) or helping her to reflect and improve their clarity about their guiding aims and values, will better enable the agent to guide her conduct autonomously. Crucially, this cannot include removing options from the set of available choices without the agent's autonomous agreement, as this would be an undue influence by our definition. At the same time, the mere provision of more options does not necessarily promote autonomy, as an agent might become confused, misled or overwhelmed, which could impair her ability to effectively govern her conduct according to her aims and values (Scheibehenne, Greifeneder, and Todd 2010; Misuraca, Teuscher, and Faraci 2016). Autonomy promotion can also take on a more visceral form. To use an example of a wearable, Cochlear implants that incorporate algorithmic technology to better provide assistance to its user is promoting autonomy by granting the user access to possibilities that would previously have been closed to them. The ways in which we can reduce autonomy can also take on a wide variety of forms. These range from the terribly overt (e.g. physical or economic coercion) to the dangerously covert (e.g. purposefully overwhelming a user with information in a user agreement so that they give in to consent fatigue and agree to it against their interest). As a rough heuristic, an intervention will be autonomy reducing if it reduces an agent's ability to achieve her rational goals.

2. Smart wearables and decision-making: relevant qualities and capacity

In this section, we explore the qualities and capacity of smart wearables that are most salient for autonomy promotion and autonomy reduction. The definition of smart wearables provided in the Introduction already presented constitutive features of smart wearables: the wearable must possess the sensors capable of collecting data, they are able to algorithmically process this data, they can employ this data in order to produce aim-guided and corrective output. The output in question can come in a variety of forms, from notifications of having achieved a fitness goal on a smartwatch, to modulating the distribution of electrical charge in a smart pacemaker. There are two differences between the sorts of outputs delivered by smart wearables and those delivered by other wearables: firstly, the smart wearable can be imparted with an aim, and will alter its output in an attempt to meet the imparted aims (within some, sometimes severe, limits), which directly leads into the

second difference, that the output in some cases can serve as a corrective to the device's previous outputs if these have missed the mark - an ability that is only possible thanks to both the features of data collection and algorithmic processing. A fitness watch, for example, can learn to give a user recommendations based on certain health aims, and then alter the output it delivers to achieve this end in the face of the data it continually collects. There are several other smart technologies that have these features, but as their name gives away, smart *wearables* represent the application of these features to worn items, items that will be physically in contact with the user's person. Indeed, any technology with these features that is physically worn by a person *is* a smart wearable.

Given this view of the wearables under discussion, what qualities do they possess that make them uniquely impactful on human decision-making? In our evaluation, wearables possess three qualities that, individually and together, work to amplify their potential impact on decision-making. We contend that it is their presence directly on a user's person - *proximity* - combined with their ease of accessibility and use - *convenience* - and omnipresence - *ubiquity* - that make them prime vectors for interventions that impact a user's autonomy. We now take a moment to discuss these qualities and why they have such significance.

To start, it might at this point be readily objected that while all wearables will be proximate by their definition, ubiquity and convenience are not necessary qualities of digital wearables. This is descriptively true, there will be wearables that prove difficult to use or that are rarely at hand. However, these wearables would precisely be deemed poor examples of this technology type. There are normative expectations regarding digital wearables - qualities that will be expected of a wearable if it is to be an excellent wearable - and convenience and ubiquity are two such expectations.⁶ As such, we should anticipate that developers and purveyors will aim to deliver these qualities, and users will expect to be able to make use of them. For this reason, we should treat these qualities as being of particular importance when we analyse the impact that these wearable technologies might have on our decision-making, though they are not necessary qualities neither are they accidental ones. It could also be objected that there are other digital technologies that are proximate, ubiquitous or convenient, and thus that there is nothing unique about wearable devices' possession of these qualities. However, though there are undoubtedly cases of overlap with other digital technologies that possess one or more of these qualities, when considering implications for decision-making, wearables occupy an arguably unique position in possessing all three qualities. Only smartphones are contenders. Indeed, the conceptual similarity to smartphones can allow us to draw some useful evidence in support of the contention that these three qualities are deeply relevant to autonomy. Studies involving smartphones have shown that if a device is in close physical proximity and readily accessible to a user it can quickly become an almost unquestioned part of a person's day to day activity and decision-making (Hamilton and Yao 2018; Reiner and Nagel 2017). The behavioural changes (and dependencies) that this possible unreflective adoption introduces are not always consciously clear to the person being impacted. And since a key benefit to wearables is precisely their convenience as ready to hand tools about which we don't have to give too much thought in

⁶ At this point it should be reiterated that our aim here is on commercial wearables, as outlined earlier. It is these wearables about which the normative expectations we have outlined hold true - not all wearables in all contexts. For example: military wearables, wearables used as part of the justice system (ankle-bracelets for tracking, for example) or wearables used in direct healthcare might not adhere to these expectations, just as they may not adhere to Autonomy Priority. There are likely different (if perhaps overlapping) normative expectations relative to each of these domains, and further work would be needed to unpack these and their implications.

day-to-day use, this is not an aspect of wearables that can simply be designed away. Nor would this be desirable, as many of the benefits of these technologies are the result of these very qualities. A further symptom of the proximity and ubiquity of wearables is that they will also collect data of a highly personal and intimate nature, including health, movement, and location data. This makes possible all sorts of useful interventions for improving an agent's life or even promoting their autonomy but raises concerns about possible infringements on privacy and the prospect of this data being misused to undermine the agent's autonomy.

In addition, though not wholly unrelated, to these qualities is the capacity possessed by wearables to *facilitate cognitive offloading by the user*. It is this capacity that is perhaps most characteristically associated with smart wearables, and what distinguishes the role that they can serve from that of non-smart wearables. By *cognitive offloading* is meant the delegation of control over the performance of a cognitive task or over the making of a decision to some device or system. This capacity is enabled through the wearable being able to collect and algorithmically process data in order to produce aim-guided and (at times) corrective output. This ceding of control can come by degrees (Dunn and Risko 2016; Risko and Gilbert 2016; Heersmink and Carter 2017). It may seem initially counter-intuitive, but this relinquishment of control can be part of a global promotion of autonomy (Kohler et al 2014; Carter 2018). The best analogy is with so-called skilled action: if I am a tennis player, I work to make certain aspects of my game habitual and reflexive - ceding control over those aspects to my automatic responses in order to be better able to accomplish my overarching aims, i.e. winning tennis matches. Similarly, we sometimes cede control to other agents in such a way that they can support our autonomy, and with the result that we are better able to respond to the reasons that enable us to best achieve our aims: think of a fitness instructor who modifies our behaviour toward our desired end frequently by usurping control over aspects of our workout regimen. Where wearables are concerned, this allows us to delegate calculating an optimum workout plan to a fitness wearable or the need to search for the nearest Indian restaurant to a smartwatch. It is important to stress that wearable technologies are able to achieve this cognitive offloading only through being able to collect and process relevant data.

3. Opportunities and concerns

But what does this mean for decisional autonomy? As is likely already clear, the combination of the three qualities and the capacity for offloading means that wearables are a double-edged sword in terms of its impact on the autonomy of our decision-making. This confluence gives rise to ample opportunities for our benefit, not least of which is the potential to promote and scaffold our autonomy. On the other hand, this selfsame combination allows wearables to act as particularly effective vectors for interventions that can reduce our autonomy. We will expand on each of the proverbial sword's edges in turn.

Wearables are able to promote decisional autonomy in four general ways: (i) the freeing up of cognitive capacity, (ii) the provision of informational input, (iii) extending the range of agency, and (iv) nudging us toward an authentic aim of ours. The freeing up of cognitive capacity is fairly self-explanatory: by facilitating cognitive offloading the wearable allows the user to focus on pursuits that they deem to be more valuable, increasing the possibility that the user is able to recognise reasons that would have otherwise been overlooked and is able to achieve her aims. (ii) has already been mentioned in passing, but the provision of otherwise unavailable information can also better allow a user to recognise salient reasons.

By “extending the range of agency” what we mean here is that the wearable makes directly possible options that were previously unavailable. Most of the obvious examples of this are where wearables are used to assist those dealing with reduced autonomy. A good example of this is the case of Simon Wheatcroft, a long-distance jogger who happens to be blind. Using a wearable device designed by the company WearWorks, which is connected to a GPS system and has several worn sensors that collect movement and proximity data that is then processed in order to provide Simon (or any user) with guidance through haptic cues. In the words of Wheatcroft, “As a blind person, you always strive for independence. But it’s a bit of a contradiction, because oftentimes, you’re using somebody with sight to become independent. What we’re trying to do is use this technology to really achieve true independence” (Sisson 2017).

A discussion of (iv) first requires some clarity on what is meant by *nudging*. To nudge an agent X as regard some decision Y, is to make changes to X’s choice architecture relevant to Y such that some preferred choice is promoted without either removing any options from the table or introducing new economic incentives (Thaler and Sunstein 2009; Felsen, Castelo, and Reiner 2013; Moles 2015; Levy 2017). The idea behind a nudge is that the agent retains full autonomy in her decision-making, but it increases the likelihood that the agent selects the choice the nudger wants. Nudges can be, and regularly are, employed in order to promote welfare or even to support the autonomy of the nudgee. Such nudges can be particularly effective when applied by wearables, thanks to the qualities of proximity and ubiquity. A smartwatch that tracks a user’s fitness data while out on a jog and then uses this data to suggest when the user should take a break is a simple but effective example of a wearable employing a nudge that can prevent an agent from overexerting themselves or aggravating a medical condition. This is clearly a case where the nudge serves to promote welfare, while at the same time not obviously infringing on the autonomy of the user. But nudges can be better than autonomy-neutral, and in some cases can actively strengthen an agent’s capacity for self-governance (Levy 2017; Niker forthcoming). Envisage the following scenario: a smoker seeks to break her addiction, and to this end she purchases a health device that can remind a user of the dangers of smoking, perhaps accompanying the warning with off-putting images, when it detects the user is smoking. The device is serving to support the agent’s autonomy by supporting their attempt to quit smoking. These nudges can also be far subtler than direct communication with the user: the layout of a user interface - colouring one option brightly while leaving the other dull, placing some qualities very visibly while placing others behind menus, etc. - can nudge users toward some choices over others. For this reason, the design of a user interface must be carefully considered, both to avoid unintended nudges and where possible to employ nudges that best support the autonomy and welfare of the user.

When we consider the challenges to decisional autonomy, there are three general categories: (i) the risk of overchoice, (ii) the risk of de-skilling and dependency, and (iii) the possibility for sludging and overnudging. The first of these has been mentioned already and refers to the - now well studied - situation where the provision of increased options serves to reduce the agent’s ability to choose the option that is in fact the best fit for her authentic aims. It is vital, therefore, that wearables should strive to provide palettes of *relevant* options in a *usable* way, and user agreements (a common environment for overchoice) should be aimed more at explainability and usability than sheer transparency or providing maximum details. This agentially useful epistemic accessibility will often only be accomplished through active user testing and feedback - the best way to know what option or information will be relevant and useful to a user is to listen to them. But it should also not be expected that users will have uniform needs in this regard. Different options and different information will be variably

useful to different agents, a commonsense fact but no less important for being so. Fortunately, where the provision of options is concerned, smart technology is well-positioned to tailor options and informational input to the needs of individual users in a dynamic fashion. However, before such tailoring can be undertaken, the user must first hurdle what can often be difficult to penetrate user agreements, some of which employ overly jargon-riddled technical or legal language that form a barrier to comprehension. Given that there are understandable legal concerns on the part of the developers and purveyors of these technologies, the legal details of such an agreement cannot be bypassed nor left up to the discretion of the user. Thus, though there is a moral reason for the drafters of these agreements to strive for explainability, this will understandably not be their only concern, and there is a moral onus on the user to educate themselves on the legal details of the agreement they enter into.

Turning to (ii), although it is true that cognitive offloading can result in a promotion of autonomy, the opposite outcome is also possible. One way in which this can occur is if a user becomes too dependent on a device, such that their own skills and decision-making ability atrophy to the point where autonomy is threatened (this is often referred to as “de-skilling” (Vallor 2015)⁷). This is most likely to occur in situations where the use of the technology becomes unreflective or habitual, precisely the danger raised by the proximate, ubiquitous, and convenient nature of wearable technologies. There are two ways in which this sort of atrophication can prove dangerous to autonomy: a) where a dependency forms on an unreliable technology and b) where the dependency stunts the development of capacities necessary for autonomy. If the technology is unreliable, then the delegation of control from the user to the technology in order to grant them greater overall control backfires. The user, if they are dependent - that is, the skill necessary to fulfill the task the technology now fulfills has atrophied away - will be left with reduced overall autonomy if the technology fails. This is not a polemic against any dependency resulting from cognitive offloading: dependency on navigation technology is a boon to the autonomy of many, and since these systems are sufficiently reliable (most of the time!) we can judge that they are autonomy-promoting. The vital takeaway from (a) is that the developers of technologies that can facilitate cognitive offloading must be thorough in assessing whether or not they are likely to result in dependencies, and if they are it is vital that the reliability of the technology be of the highest order. However, even if reliability is not a concern, there are still some dependencies that may be pernicious to decisional autonomy. If the dependency results, directly or indirectly, in the stunting or loss of a capacity necessary for decision-making, then we have *pro tanto* moral reason to oppose the dependency-inducing technology, one that will rarely if ever be overridden in the commercial realm.

Lastly, we have (iii). Following current convention, we call nudges that nudge an agent against their interests: *sludges* (Thaler 2018). Given our account of autonomy, such sludges are autonomy reducing as they work against an agent achieving their aims. These sorts of interventions can take many shapes but are usually employed with the interest of increasing profits at the user’s expense. Placing expensive but unhealthy candy at the exit of a

⁷ Vallor concentrates here attention on moral de-skilling, where the atrophied skill is one pertaining to moral conduct or judgement, but the point generalises equally to all skills that (i) can be offloaded and (ii) are capable of atrophy. It is also important to note that it is not only the atrophication of moral skills that have moral significance, as the loss of certain non-moral skills can result in otherwise avoidable harms - consider the loss of attentional awareness (not a strictly moral skill), which can result in lowered life success for those afflicted (Baumgartner et al 2018). For a discussion of some of the impacts these technologies might have on our cognition, see (Firth et al. 2019).

supermarket is an example of a measure that can function as a sludge, inducing customers to purchase these products even when it works against their authentic aims. Combating sludges is often best achieved by informing users about their presence and the danger they pose. Awareness of a nudge or sludge, though not foolproof, can go a long way to helping people resist its possible effects on their decision-making.

Apart from sludges, there are two other ways in which nudges can undermine autonomy. Firstly, our aims and values can often prove very endogenous, leaving us vulnerable to being nudged away from our own authentic self-government. This is particularly true if nudges operate by bypassing our deliberative capacities (Grüne-Yanoff 2012). Secondly, nudging can serve to prevent or impair the development of capacities necessary for autonomy by cutting an agent off from irreplaceable learning experiences (Blöser et al 2010). As well described in Niker et al. (forthcoming):

Consider an adolescent who, rather than finding their own way in the world by ‘learning from their mistakes’, has parents who remove obstacles from their path...[i]n essence, these adolescents live in an environment that is designed by choice architects (their parents, in this case) to make the best decisions most likely. They may end up with decisions that are welfare-promoting, or even ideal in some sense, but there is less opportunity for them to develop the fundamental skills involved in decision making. The worry is that a similar sort of diminishment of human decision-making competencies is going on in a world structured by public policy nudges.

This can occur even when the nudger has the best intentions and is exacerbated when the nudgee is the target of many concerted nudges or the source of the nudging is unreflectively integrated into the nudgee’s decision-making. Wearables run the risk of bringing about precisely this through their proximity, ubiquity, and convenience. One of the best and simplest ways to combat this risk is to inform users about how they are being nudged - or will be nudged. This will likely reduce the efficacy of at least some nudges, which often work best when undetected, but this is a price that should be paid in seeking out the appropriate balance, especially in a commercial context.⁸

Concluding remarks

In this work we have investigated the opportunities and challenges for decision-making raised by smart wearable technologies. First, we unpacked the moral importance of decision-making, showing that the moral value present was grounded by considerations of autonomy. In this first section we also introduced the important idea that different external influences on an agent can be autonomy promoting or autonomy-reducing, and that there is at least a *pro tanto* moral obligation in favour of promoting autonomy in ourselves and others. In Section 2, we discussed the qualities and capacity of wearables that are likely to have the most significant and most novel impact on our decision-making abilities: the qualities of

⁸ One excellent strategy to consider then undertaking nudging is that raised by Levy (2017) and discussed by Niker et al. (forthcoming): that wherever possible so-called nudges to reason can be employed in order to foster increased reasons-responsiveness and thus promote autonomy. The key with these nudges is that they actively seek not to promote a certain action, but rather to promote active reflection on the part of the nudgee by engaging filters that are themselves “partly constitutive of our reasoning” itself (Levy 2017). If the nudge is such that its effect on the nudgee is via a process that forms a part of the normal functioning of our reasons-responsive mechanisms, then it does not run afoul of our definition of undue external influence from Section 1 and so does not threaten autonomy.

The Moral Opportunities and Perils of Smart Wearables for Decisional Autonomy. By Niël Henk Conradie. Research project: InviDas. November 2021. <https://invidas.gi.de/>

ubiquity, proximity, and convenience, and the capacity to facilitate cognitive offloading. Thirdly and finally, we identified and considered the primary moral opportunities and concerns presented by smart wearables for human decision-making and put forward a few suggestions as to how the former can be supported and the latter avoided.

List of References

- Baumgartner SE, van der Schuur WA, Lemmens JA, and te Poel F (2018). The Relationship Between Media Multitasking and Attention Problems in Adolescents: Results of Two Longitudinal Studies. *Human Communication Research*, 44(1), 3-30.
- Blöser C, Schöpf A, and Willaschek M (2010). Autonomy, experience, and reflection. On a neglected aspect of personal autonomy. *Ethical Theory and Moral Practice*, 13(3), 239–253.
- Carter JA (2018). Virtue Epistemology, Enhancement, and Control. *Metaphilosophy*, 49(3), 283-304.
- Christman J and Anderson J (Eds.). (2005). *Autonomy and the challenges to liberalism: New essays*. Cambridge: Cambridge University Press.
- Darwall S (2006). *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge: Harvard University Press.
- Dian JF, Vahidnia R, and Rahmati A (2020). Wearables and the Internet of Things (IoT), Applications, Opportunities, and Challenges: A Survey. *IEEE Access*, (8), 69200-69211.
- Dunn T and Risko EF. (2016). Toward a Metacognitive Account of Cognitive Offloading. *Psychology, Medicine, Computer Science*, 40(5), 1080-1127.
- Felsen G, Castelo N, and Reiner PB (2013). Decisional enhancement and autonomy: Public attitudes towards overt and covert nudges. *Judgment and Decision Making*, 8(3), 202–213.
- Firth J, Torous J, Stubbs B, Firth JA, Steiner GZ, Smith L, Alvarez-Jimenez M, Gleeson J, Vancampfort D, Armitage CJ, and Sarris J (2019). The “online brain”: how the Internet may be changing our cognition. *World Psychiatry*, 18, 119-129.
- Frankfurt H (1998). Autonomy, necessity, and love. In: *Necessity, volition, and love*. Cambridge University Press, Cambridge, 129–141
- Grune-Yanoff T (2012). Old wine in new casks: libertarian paternalism still violates liberal principles. *Social Choice and Welfare*, 38, 635–645.
- Hamilton KA and Yao MZ (2018). Cognitive Offloading and the Extended Digital Self. In Kurosu M (Ed.), *Human-Computer Interaction. Theories, Methods, and Human Issues. HCI 2018. Lecture Notes in Computer Science, vol 10901*. Springer, Cham.
- Heersmink R and Carter JA (2020). The philosophy of memory technologies: Metaphysics, knowledge, and values. *Memory Studies*, 13(4), 416-433.
- Köhler S, Roughley N, and Sauer H (2017). Technologically blurred accountability? Technology, responsibility gaps and the robustness of our everyday conceptual scheme. In Ulbert C, Finkenbusch P, Sondermann E, and Debiel T (Eds.), *Moral Agency and the Politics of Responsibility* (pp. 51-68). London: Routledge.
- Korsgaard CM (1996). *The Sources of Normativity*. Cambridge University Press.
- Levy N (2017) Nudges in a post-truth world. *Journal of Medical Ethics*, 43, 495–500.

The Moral Opportunities and Perils of Smart Wearables for Decisional Autonomy. By Niël Henk Conradie. Research project: InviDas. November 2021. <https://invidas.gi.de/>

Misuraca R, Teuscher U, and Faraci P (2015). Is more choice always worse? Age differences in the overchoice effect. *Journal of Cognitive Psychology*, 28(2), 242-255.

Moles A (2015). Nudging for Liberals. *Social Theory and Practice*, 41(4), 644-667.

Niker F, Felsen G, Nagel SK, and Reiner PB. (Forthcoming). Autonomy, Evidence-Responsiveness, and the Ethics of Influence. In Blitz M and Bublitz JC (Eds.), *Neuroscience and the Future of Freedom of Thought*. Hampshire: Palgrave-Macmillan.

Reiner PB and Nagel SK (2017). Technologies of the Extended Mind: Defining the Issues. In Illes J & Hossain S (Eds.), *Neuroethics: Anticipating the Future* (pp. 108-122). Oxford Scholarship Online.

Risko EF and Gilbert SJ (2016). Cognitive Offloading. *Trends in Cognitive Sciences*, 20(9), 676-688.

Scheibehenne B and Greifeneder R (2010). Can There Ever be Too Many Options? A Meta-analytic Review of Choice Overload. *Journal of Consumer Research*, 37, 409-425.

Seneviratne S, Nguyen T, and Lan G (2017). A Survey of Wearable Devices and Challenges. *IEEE Communications Surveys & Tutorials*. DOI: 10.1109/COMST.2017.2731979.

Sisson P (2017). Beyond the finish line: how technology helped a blind athlete run free at the New York Marathon. *The Verge*, Nov 6. Online: <https://www.theverge.com/2017/11/6/16610728/2017-new-york-marathon-blind-runner-wearworks-wayband-simon-wheatcroft>

Thaler RH (2018). Nudge, not sludge. *Science*, 361(6401), 431.

Thaler RH and Sunstein CR (2009). *Nudge: Improving decisions about health, wealth, and happiness*. New Haven: Yale University Press.

Vallor S (2015). Moral Deskillling and Upskilling in a New Machine Age: Reflections on the Ambiguous Future of Character. *Philosophy & Technology*, 28, 107-124.

Westlund AC (2009). Rethinking Relational Autonomy. *Hypatia*, 24(4), 26-49.