



Doing Away with the Agential Bias: Agency and Patency in Health Monitoring Applications

Nils-Frederic Wagner¹

Received: 14 July 2017 / Accepted: 9 April 2018 / Published online: 23 April 2018
© Springer Science+Business Media B.V., part of Springer Nature 2018

Abstract Mobile health devices pose novel questions at the intersection of philosophy and technology. Many such applications not only collect sensitive data, but also aim at persuading users to change their lifestyle for the better. A major concern is that persuasion is paternalistic as it intentionally aims at changing the agent's actions, chipping away at their autonomy. This worry roots in the philosophical conviction that perhaps the most salient feature of living autonomous lives is displayed via agency as opposed to patency—our lives go well in virtue of what we do, rather than what happens to us. Being persuaded by a device telling us how to conduct our lives seemingly renders the agent passive, an inert recipient of technological commands. This agential bias, however, has led to a marginalization of patential characteristics that are just as much part of our lives as are agential characteristics. To appreciate the inherent interlocking of acting and being acted upon, it is vital to acknowledge that agency and patency are correlates, not mutually exclusive opposites. Furthermore, it is unclear whether an action can only count as agential so long as its causes are internal. Drawing on the extended mind and extended will framework, I argue that mHealth applications merely serve as volitional aids to the agent's internal cognition. Autonomously set goals can be achieved more effectively via technology. To be persuaded by an mHealth device does not mainly—let alone exclusively—emphasize patency; on the contrary, it can be an effective tool for technologically enhancing agency.

Keywords Mobile health monitoring · Smartphone applications · Persuasion · Autonomy · Paternalism · Agency and patency

✉ Nils-Frederic Wagner
nils-frederic.wagner@uni-due.de

¹ Department of Philosophy, University of Duisburg-Essen, Room: LE 329, Forsthausweg 2, 47057 Duisburg, Germany

1 Introduction

It is estimated that by 2020, roughly 70% of the world's population will be using smartphones; currently, the worldwide dissemination is at around 50%.¹ While such devices continue to be used mostly recreationally, the burgeoning tendency to *smarten up* our lives has long reached the medical sphere. New technological means enable an ever-increasing number of users to track and analyze a vast amount of sensitive, personal health-related data. Mobile health applications (commonly referred to as “mHealth apps”) that typically run on smartphones are the most pervasive form of such devices. A recent study across 16 countries found that, as of today, around one third of the online population uses mobile health devices.² Core features of most such apps already at hand are the tracking of movements, nutrition, and sports activities. Physiological parameters like heart rate and blood pressure that are known to highly correlate with emotional states, and wellness factors such as quality of sleep, and social interaction are also monitored but may require gadgets like chest straps or heart rate watches. These smart devices facilitate a previously unheard-of efficacy of self-monitoring. With the rise of such novel, intimate technologies, a variety of philosophical issues crop up concerning, most pertinently, data security, responsibility, paternalism, autonomy (Krieger 2013; Owens and Cribb, *forthcoming*), as well as conflicts of interest between different stakeholders.

Concerns about users' autonomy become ever more pressing since a growing number of such applications do not merely collect data, but also aim at persuading users to change their lifestyle for the better, i.e., living a healthier, more active life. To achieve this goal, app designers utilize a variety of persuasive strategies that potentially erode users' autonomy and threaten their agency.

Here's a case that shall help illustrating the concern. Suppose you are going for lunch with your colleague. When you are about to order, a message from your recently acquired mHealth app pops up telling you to have the healthy salad option (needless to say, you would much rather have fish and chips). Having salad, the device says, will lower your cholesterol level which in turn will make you feel better in the long run. From automatically checking the weekly cafeteria menu online, the device knows the food options and comes up with the healthiest choice for everyday. This is done according to a complex algorithm, taking into account both your physical parameters as well as all previously chosen meals ever since you have been using the device to ensure that your diet is more balanced. You are in a clear epistemic disadvantage here; matters that are considered by the app to generate choices are just too complex for you to fully comprehend. That way, the app has an expertise you lack. Let us suppose further that if you go for salad, the device rewards you by allocating health points to your account, say in the form of green leaves. Since this app is quite popular among your friends, you have entered into a competition. Whoever has collected most green leaves by the end of the year will be declared winner and can look forward to being invited to a fancy getaway by the other participants. Lots of reasons to go for salad, it seems. But what if you cannot resist your cravings and go for fish and chips anyway? In

¹ <https://techcrunch.com/2015/06/02/6-1b-smartphone-users-globally-by-2020-overtaking-basic-fixed-phone-subscriptions/> (retrieved May 2017).

² <http://www.gfk.com/global-studies/global-studies-fitness-tracking/> (retrieved June 2017).

that case, the app will deduct the green leaves you earned yesterday and will also send a message to your friends letting them know that you have indulged in some tasty but unhealthy meal.

When imagining such cases, most of us, I take it, feel a certain unease regarding our autonomy. Is it really you deciding to have salad for lunch? Or is it the app deciding for you? Since you are at an epistemic disadvantage, the app knows better what is good for you anyway. So, on the face of it, it looks as though you have been paternalized by some technological device, perhaps a non-human agent, using more or less subtle elements of persuasion (maybe even manipulation) to impose its will on you. Also, it employs motivational triggers by introducing elements of reward and competition to nudge you into complying.

In what follows, I focus on whether persuasive mHealth apps do indeed dislodge autonomy by constituting a paternalistic intervention into people's lives. I argue that, despite appearances, there are good reasons to believe that these systems do not per se pose a threat to agency. Under certain conditions, mHealth apps even bear the potential to technologically ameliorate agency.

I proceed as follows: Firstly, I present both paradigmatic views that argue for mHealth apps' potential to enhance users' autonomy and paradigmatic views that argue for mHealth apps' potential to threaten users' autonomy. I then sort out the conceptual interrelation of autonomy, persuasion, and paternalism in the context of mHealth apps. Subsequently, I make some remarks on how these concepts figure in the common understanding of agency and patency in philosophy of mind and action. I argue that a widespread agential bias has led to an underappreciation of patential concerns that make up significant proportions of our lives. Finally, I make the case for understanding some persuasive elements of mHealth apps as what I shall call "volitional aids" which are, considering the theories of extended mind and extended will, part of the agent's own, albeit extended cognitive architecture as opposed to external interferences, suggesting that, understood in this way, some of these apps are effectively outsourced parts of our minds and wills that potentially enhance agency.

2 Autonomy, Persuasion, Paternalism

Owens and Cribb are among those who have recently argued for the autonomy enhancement potential of mHealth apps. In particular, they think that such apps can foster users' deliberation and decision-making capacities: "By providing access to biomedical data and generating awareness of habits, behaviours and performances, there is good reason to think these technologies can support processes of deliberation about health that enhance their users' procedural autonomy. For example, information about one's heart rate, sleeping patterns, mobility or calorific intake might help people make important decisions that directly affect their health" (Owens and Cribb, [forthcoming](#), 5).

Recent empirical studies indicate that some users feel autonomous and motivated when employing mHealth apps in their daily routine. It remains elusive how much theoretical weight should be given to a sample of users' assessment of such apps. Nonetheless, as the following summary of user reports illustrates, there is reason to at least surmise some autonomy enhancement potential: "Users reflected positively on the

use of the apps, with one user felt that the autonomy-supportive style was evident in terminology used. Users felt motivational value from seeing steps, styles and advice. User attitudes reinforced autonomy stating it made use of the device more engaging and positively influenced sustained or repeat use. Generally, users enjoyed the level of autonomy they were granted by the apps. However, some stated a need for apps to balance autonomy with more self-directed goal creation to support their engagement” (Asimakopoulos et al. 2017, 7).

In contrast to the views such sketched, several scholars have argued that technology in health care in general, and mHealth apps in particular, pose threats to users’ autonomy. In what follows, I highlight some of the most pressing worries such views articulate before relating these views to an analysis of persuasion and paternalism in this context.

Timmer et al. argue that due to new technological means of persuasion it “might be harder for the individual to make an autonomous choice about the goals he is being persuaded to, or whether he consent[s] to the use [of] persuasive technologies” (Timmer et al. 2015, 196). They argue that safeguarding autonomy in these new means of persuasive technology is even more important when the setting of their application takes place in sensitive contexts like health care (ibid., 197). When persuasive technologies appear in what the authors call “collective applications,” “for instance in healthcare and insurance—research is needed on the role of these third parties as providers of persuasion and how they impact the users’ autonomy” (ibid., 201).

Several authors that critically engage with new mHealth technologies such as Lanzing point to a tension between disclosing sensitive personal information and safeguarding one’s autonomy: “self-tracking breaks down informational privacy boundaries that otherwise enable autonomous self-presentation within different social contexts” (Lanzing 2016, 10). Lanzing further thinks that users’ autonomy is comprised because of a potential breach of information privacy in mHealth apps, since, one the one hand, users are encouraged to collect and share as much data as possible, both to increase functionality and persuasion. But on the other hand, privacy of information is a hallmark of living autonomously. Success stories about empowerment, self-control, and self-improvement camouflage the reality of decontextualization, thinks Lanzing, where we expose too much to an undefined (future) audience, which limits our capacity to run our lives for ourselves. Altogether, this constitutes a violation of users’ privacy that can undermine their autonomy on a more fundamental level (ibid., 15).

Nordgren also places particular importance on privacy in personal health monitoring, submitting that frequently “the user has no autonomy regarding which information is to be collected, transmitted, processed and used” (Nordgren 2015, 155). To avert this issue, Nordgren suggests a context-sensitive balancing of automated privacy protection that might be feasible in some circumstances and autonomously chosen privacy protection that might be called for in other circumstances (ibid., 163).

Sharon holds against the idea of mHealth apps as empowering users that, “self-tracking for health is disempowering, insofar as it invites an increased control of others—health promoters, friends and followers, and even the internalized health promoter of one’s own super ego—over oneself” (Sharon 2017, 99). She further posits that “discourses of empowerment and healthy citizenship are seen as concealing economic realities that are often detached from the interests of citizens and patients

and of creating new forms of discipline, subjection, and social control—of imposing limits on the autonomy of individuals” (ibid., 106f.).

Even though these views focus on different aspects of autonomy, they commonly see threats to users’ autonomy as a problematic interference with their agency that ought to be circumvented since agency is something worth aspiring to when it comes to living well. For this reason, I describe such views as having an “agential bias.” In Section 3, I say more about how this agential bias might be rooted in the common tendency in philosophy to see our lives as going well, first and foremost, in virtue of agential features.

Although most of the literature leans towards either of the two sides just sketched, some authors see both negative and positive aspects of personal health monitoring regarding users’ autonomy. Here is one such view: “However, PHM [personal health monitoring] can also restrict the lifeworld, impinging the system’s economic and power concerns on the individual lifeworld such that restrictions are placed or information demanded in order to maintain institutional structures. Hence PHM has the potential to act both as the repressive father, dictating behaviour and routine and demanding information for his own purposes, or the supportive mother offering both reassurance but also an environment which supports the autonomy of the patient” (Mittelstadt et al. 2014, 50).

2.1 Key Aspects of Persuasive mHealth Devices

The semantics of persuasion plays a central role in assessing the issue as to whether its application in mHealth apps poses a threat to users’ autonomy, possibly eroding their agency. To a first approximation, persuading someone to doing (or omitting)³ something is to intentionally try changing their actions via their convictions and intentions that lead to action. Whereas persuasion has a largely positive connotation in social psychology and health science (Cialdini et al. 2005), its reputation in the philosophy of action is rather seedy; sometimes, persuasion is seen as akin to (or at best in between) manipulation and convincing (O’Keefe 2012).

Persuasive technologies are generally designed such that they provide technological means to intentionally, and often permanently, change users’ behavior via their convictions and intentional states by constantly providing feedback on what is understood as inadequate behavior and by incentivising in various ways what is deemed desired behavior (cf. Fogg 2003).⁴ Persuasion is thus inherently normative as its main rationale is not to merely describe something or inform someone, but to persuade, or as some argue, to manipulate people into doing something.

Key characteristics of persuasive technologies in the context of mHealth applications are that they work with body sensors, implemented on smart devices such as smartphones and smart watches that continuously collect and display the recorded data, that they provide real-time persuasive feedback, that they function widely automated without the need for human control, that they are customized so as to

³ Whenever I talk about “doing” in the context of persuasion and autonomy, I take “omitting” to be implied. For brevity’s sake, I hereafter omit “omitting.”

⁴ Davis et al. ’ (2015) scoping review surveys health-related behavior change theories as they are put forward in the social sciences.

accommodate users' specific needs, and that they are context-sensitive, taking into account users' current condition (Koelle et al. 2014). Furthermore, the design of such applications is supposed to carefully consider technological, social, and interactive components of human-computer interfaces that enable a smooth human-computer interaction.

Even though persuasion is generally taken to be a sustained effort to change someone's behavior, with paternalism and patronization lurking in the shadows, ideally, there are various measures in place that prevent technological persuasion from being unethical: (1) the desired behavior change should be achieved without deception (i.e., neither in terms of concealing the striven-for outcome nor in terms of disguising the measures taken to achieve that goal), (2) users should voluntarily decide to use such technologies, and (3) the intended goals of persuasion should be kept transparent (Chatterjee and Price 2009), as much as this is possible without running into problems of *persuasive backfiring* (i.e., the triggering of unintended outcomes of behavior change).

Before discussing principles of ethically sound persuasion in more detail, I turn to motivate potential threats of paternalistic interventions to people's autonomy. Firstly, I discuss defining conditions and standard cases of paternalism and then relate main elements of these scenarios to technologically aided paternalistic interventions.

2.2 Paternalistic Interventions and the Principle of Respecting Autonomy

The widely held principle of respecting autonomy suggests that persuasion chips away at the agent's autonomy since it figuratively (or, as the case may be, literally) talks the agent into doing something they would not have done otherwise, out of their own free will. It is frequently argued that such persuasive interventions violate the principle of respecting autonomy since they constitute a form of paternalism, interfering with the agent's own volition; albeit motivated by the conviction that the agent will be better off when persuaded into changing their behavior accordingly (Enoch 2016). Some hold that the basis for a behavioral adaptation so achieved is not that the agent was rationally convinced to do so, but rather deceptively manipulated into doing so (Spahn 2012). This line of thought stems from persuasion's rather shady reputation in philosophy, where it is often seen as inherently paternalistic and thus at odds with respecting the agent's autonomy.

In numerous writings, perhaps most succinctly in the *Stanford Encyclopedia of Philosophy*, Gerald Dworkin (1972, 2005, 2015, 2017) essays a definition of paternalistic interventions that helps understanding why persuasion is often held in disesteem in philosophy, and why it might be seen as posing a threat to agency.

Dworkin suggests the following conditions as an analysis of *X acts paternalistically towards Y by doing (omitting) Z*:

1. Z (or its omission) interferes with the liberty or autonomy of Y.
2. X does so without the consent of Y.
3. X does so only because X believes Z will improve the welfare of Y (where this includes preventing his welfare from diminishing) or in some way promote the interests, values, or good of Y.

Even though Dworkin does not explicitly say that X and Y are two distinct agents, each having their own set of intentions, his analysis implicitly suggests that that is what he has in mind. Paternalism, then, occurs when one agent imposes their will on another agent with the benevolent, albeit patronizing intent of promoting (or at least preserving) the other's well-being or interests more generally.⁵

Given the involvement of two distinct agents in paternalistic interventions à la Dworkin and others, (2) in conjunction with (1) constitutes a violation of Y 's autonomy *eo ipso*. One cannot both respect someone's autonomy and, at the same time, interfere with their autonomy by acting on them without having previously obtained informed consent. However, the conjunction of (1) and (2) leaves open whether interfering with someone's autonomy were morally permissible if the interfered-with, acted-upon, or paternalized agent had given their consent.⁶ (3) subverts Y 's autonomy inasmuch as it implies that X is in a better epistemic position to judge (or is in some other way more competent) than Y themselves with regard to figuring out what actions and intentions that lead to action are conducive to Y 's well-being; thus questioning their decision-making capacity, rationality or whichever agential feature might be required for the task at hand. (1) most obviously relies on the assumption that the paternalistic action Z is performed or initiated by X , whereby X is an agent in their own right, having beliefs and intentional states they want to impose on Y . Given that (2) supposes that X performs Z on Y without their consent, evidently, X and Y are taken to be two distinct agents.⁷ Now, while this might hold true in standard cases of paternalism (for example, a parent taking away their drunken child's car keys to prevent them from causing an accident or from getting pulled over for DUI), when mHealth apps are concerned, it is far from clear whether there actually are two distinct agents at play. The question, then, becomes: are such apps "covert agents" pushing someone else's agenda by proxy with the means of technological persuasion, or are these apps, perhaps, just an extension of the agent's own volition? It goes without saying that the "covert-agent-concern" largely depends on the app-creator's intentions and on its design. For example, an app commissioned by an health insurance company with the aim of reducing costs by persuading their policyholder to, say, quit smoking, might well be paternalistic in that another agent (a team of app-designers on behalf of the company's CEO) tries imposing their will on users by means of technological persuasion. But importantly, this need not always be the case, nor is this necessarily so in mHealth apps.

Dworkin's three conditions of paternalism rely on the *prima facie* tenable assumption that X and Y are distinct agents, each representing their own set of intentions. Paternalism, then, seems to occur just in case X acts upon Y by meeting at least one of Dworkin's three conditions. While this might be reasonable in (2), since this condition requires another agent *ipso facto*, it is not necessarily true in (1) and (3). I venture that

⁵ This reading is suggested by many other accounts of paternalism. Michael Cholbi (2017), to name just one example, has recently put forward a version of rational will that ranks the wrongfulness of paternalistic interventions in terms of the extent to which such acts replace the paternalizee's practical rationality with the paternalizer's and the degree of mistrust in the paternalizee's rational agency displayed by the paternalistic intervention.

⁶ This question requires a separate treatment, but is not of central importance for my purposes.

⁷ The issue as to whether one can, in principle, interfere with one's own autonomy or whether this necessarily requires another agent is interesting, but a detailed analysis thereof must remain a task for another day. At first blush, it looks as though actively deciding to renounce one's autonomy is itself an act of autonomy.

(1) and (3) need not always involve another agent. Why is that? I can, in principle, act upon myself paternalistically for example by forming distal intentions and putting measures in place that will make my future self comply. Think of, for example, Ulysses tying himself to the mast to resist the Sirens' song, or Parfit's Russian nobleman requesting his wife to hold him to the promise to distribute large portions of his wealth once he reaches a certain age even though his older self might have a change of heart; although, this case is less clear, since the anticipated change in attitude does not need to involve a decline in rationality. Nevertheless, once the implicit claim that being acted upon inevitably requires two distinct agents is dropped, it is much more contentious whether paternalism always interferes with agents' autonomy. In fact, a situation where *X* acts upon *Y*, where *X* is not a distinct agent but a technological device acting *in the service* of *Y* (i.e., as an extension of, and not in opposition to, *Y*'s will), might not constitute a case of paternalism at all—surely, it does not seem to pose an obvious threat to the agent's autonomy. It might even be a genuine expression of the agent's autonomy.

Before returning to the idea of expressing one's autonomy through being acted upon more thoroughly, I shall address the following question: What is it with this apparent contrast of paternalism and agency? It appears that paternalistic interventions address us primarily as patients and thereby chip away at our autonomy. In what follows, I want to change what I take to be a misguided focus by suggesting that patency need not be agency-eroding but can, under certain circumstances, be a display of agency. Admittedly, expressing one's autonomy by being acted upon is unusual and difficult to grasp since autonomy is ordinarily displayed by one's actions, not by what happens to us (for lack of a proper noun; "inaction" does not seem to do the trick since we are not necessarily inactive when something happens to us). The underlying dichotomy between actions and things that happen to us plays a crucial part in agency's reputation as displaying inherently active features of agents' lives. So, it might be worth taking a closer look at the allegedly opposing concepts of agency (as in acting) and patency (as in being acted upon). Such an analysis shall help revealing that, perhaps, there is not such a sharp divide between these two aspects of people's lives after all.

I now turn to make some remarks on the common tendency in philosophy of mind and action to underappreciate patiential traits of our lives and to spell out how this agential bias might have given rise to the misconception of technological persuasion as agency-eroding. In a subsequent step, I hope to show that exercising patiential characteristics of people's lives can, under certain conditions, enhance agency—if not paradigmatically, then at least more commonly than initially thought.

3 The Agential Bias

The wide notion of "patients" and "patency" as technical terms in philosophy has a different connotation than the narrow notion of patients in ordinary language, particularly in healthcare settings. Philosophically, being a patient describes, broadly, the passivity of someone who undergoes some action or to whom something is done. This passivity of patients is chiefly contrasted with the activity displayed by agents. Roughly, the philosophical contrast between agents and patients is captured by the slogan: "agents do things, whereas things are done to patients." This philosophical dichotomy

is initially independent of the settings in which agents act and patients are acted on. The medical notion of patients, which is commonplace in ordinary language, locates patients in the vicinity of health care. Patients are thus people that suffer from a medical condition for which they receive medical treatment, either in outpatient or inpatient care. In what follows, I am mainly concerned with the philosophical notion of patency. In the context of mHealth apps, however, “philosophical patients” are in some sense also “medical patients,” but only contingently so. My arguments do not, therefore, rely on the medical notion of patients. I say more about this in Section 3.1.

Agency enjoys a considerable privilege in philosophy of mind and action, as does moral agency in ethics. Philosophers often emphasize that our lives go well in virtue of what we *do*, rather than in virtue of what *happens to* us (Lott 2016).⁸ A paradigmatic way of phrasing the agential bias is put forward by Mark LeBar when he says that his view of living a good human life “is agentist, not patientist ... we are first of all agents, who live by acting on their world” (LeBar 2013, 69 f.). To be an agent is, by and large, to actively partake in life; agents mold the world around themselves. Patients, on the other hand, are people to whom things happen; passive sufferers, molded by life’s happenings. On that view, when our lives begin, we start out as depended patients, and, if everything goes well, in the course of our adult lives, we evolve into fairly independent agents. It goes without saying that this is but an ideal we aspire to, never to be fully achieved.

To appreciate the conjunction of acting and being acted upon, it is important to acknowledge that agency and patency are *correlates*, not mutually exclusive opposites. Soran Reader (2007) neatly describes how the active features of agents’ lives that are taken to be exclusive to agency (on her account action, capability, choice, and independence) all have a corresponding “other side” to them (on her account, passion, liability, necessity, and dependency, respectively). Reader characterizes this other side of agency as “a complementary aspect which necessarily accompanies the aspect valorised as ‘positive’ and assumed to furnish the essence” (ibid., 588) of what makes an agent.

Focusing on action, Reader carves out two aspects in which agency is necessarily accompanied by its complementary other side: Firstly, she claims that agents themselves suffer from their action, and thus always are, inevitably, in some relevant sense, patients as well as agents. For example, when I ride my bike, pedaling hard, I do not just move my bike forward, but I also suffer the pedals’ resistance. Another example Reader cites is that when I hit you, it is not just you that suffers from my punch, but it is also I that suffers from your resistance to the blow. In the second sense, according to Reader, every action requires a patient at the receiving end of that action. The person being hit in the previous example is a most obvious case of a patient. So, every action requires both an agent initiating the action (whereby the agent is to some extent also a patient with respect to that very action) and a patient being passively affected by the action. As we have seen, in some cases, agent and patient involved in a particular

⁸ There is much to be said about the interplay between acting and being acted upon. Mikael M. Karlsson (2002), for example, cashes out the distinction between things we do and things that happen to us in an Aristotelian attempt of “self-movement.” Richard Taylor (1982) questions the metaphysical distinctiveness of action and suggests a more practical approach, claiming that we decide whether something is an action or someone an agent when we encounter them; as we go along, so to speak.

situation can be one and the same person. When I lift my cup of tea, I am both an agent sipping from my cup and a patient suffering the cup's touch at my lips.⁹

Since agency and patiency are complementary parts of a person's life, Reader submits that ascribing agency metaphysical primacy in the constitution of personhood is unfounded. It is up for debate whether one must follow Reader all the way to this metaphysical conclusion. Certainly, as an anonymous referee rightly pointed out, additional arguments are needed to confute the metaphysical claim Reader prematurely rejects; perhaps agency does deserve priority in the metaphysical constitution of personhood. All the same, a decisive decision on this matter is not necessary for my purposes, and so I remain agnostic about the matter here. From a more practical point of view, drawing attention to the tight linkage between agency and patiency is a valuable insight that should help alleviate the agential bias by emphasizing the complementary nature of these two integral aspects of people's lives.

Another reason why it is difficult to do away with the agential bias is the common misconception to view patients on a par with mere *objects*, forfeiting or lacking agential features altogether. Along these lines, Krakauer issues a Heideggerian and Foucault inspired worry regarding the technologically induced exposure of agents as mere objects in healthcare technology: "This challenging or provoking of beings to expose themselves as objects which thereby also poses or establishes beings as objects is precisely what Heidegger calls 'the essence of technology.' Foucault took as his task to follow the path indicated by this Heideggerian thought through the language of medicine. Foucault's labor of listening to medical language hears that even the autonomous individual, the subject itself, has acquired, and been reduced to, 'the status of an object'" (Krakauer 1998, 533). But this is not so. When I am being acted on, even in the crudest sense, I thereby do not cease to be an agent altogether. It is just that I might not currently exercise most of my agential features. Now, the idea is to remedy the agential bias that marginalizes patiential parts of people's lives by acknowledging that agency is manifested not only when we are acting as agents but also when we are being acted upon as patients, and furthermore, to see that we are, inevitably, patients all the time. Since, as mentioned previously, every action has both agential and patiential characteristics. And so, these patiential parts are no failure, nor are they of lesser value than agential characteristics in attaining agency. Patiency dialectically completes agency. In seeing that agency presupposes patiency, we might find reason to drop the idealization of agential characteristics as the main components for valuing our lives in favor of a more balanced view that can help appreciating patiential characteristics just as much. What happens to us as patients, in acting and in being acted on, may define us

⁹ An anonymous referee pressed the point that the interpretation of predicates like "suffering" seems very generous here. What is meant to be shown by emphasizing things like suffering (and similar patiential expressions) is the idea that in every instance of human action, there is always simultaneously an aspect in which things passively happen either to someone else involved in that action or to that very agent initiating the action (this is why I have earlier described agency and patiency as correlates). The reviewer further says, and rightly so, that it would be more intuitive to say that agents sometimes also "overcome," for example, some "resistance" but that this does not make them patients. This might be so, but I think there is still reason to resist this intuitive appeal and instead hold on to the conceptual conviction that agents are in some relevant sense always also passive with regard to their actions. "Overcoming" some resistance, in this example, suggests that the agent actively does so overcome, when in fact they might just passively "endure" something that happens to them, something that is an integral part of the action but that the agent has nonetheless no influence or control over and is thus patiential with regard to the action.

as much as what we do as agents. And so, being a patient all the time is not as such a reduced or unpleasant condition but an unavoidable fact about us. In order to appreciate the complementary nature of agency and patency, we ought to explore what passively happens to agents, what constraints them, the contingencies they are subjected to, as well as their display of active, agential characteristics.

3.1 Patency and Paternalism in Health Care

When it comes to medical settings, the notion of patency in relation to paternalism requires an additional treatment from what I have said so far concerning the broader philosophical notion of patients and paternalism. The special relation between medical patients and healthcare professionals is constituted by two assumptions. For one, medical patients rely and trust on the expertise of healthcare professionals and base their decisions on what clinicians recommend. Healthcare professionals, on the other hand, are supposed to act in ways that are both beneficial to the patients in their charge and at the same time respect patients' autonomy with regard to their medical care. When healthcare professionals act on the basis of what is good for their patients alone, and thereby ignore patients' autonomy, they act paternalistically. The asymmetry of expertise between patients and healthcare professionals allows for asking when, if ever, medical paternalism is called for. I cannot attempt to give a decisive answer to this question here, but I think that Groll (2014) has a point when he argues that the burden of proof should lie on those who think that medical paternalism is sometimes justified.

Since mHealth apps are not only used by healthy individuals to track their fitness level but also prescribed by physicians to monitor medical conditions, the distinction between the philosophical and the medical notion of patients becomes blurry and so does the question as to whether a medical form of paternalism, which potentially involves (technological) persuasion, might be acceptable in such cases. A further complication that makes a decision as to if and when medical paternalism in the context of mHealth can be appropriate is the fact that users differ vastly in their health literacy which impacts users' autonomy. What could be an enhancement of someone's autonomy, might constitute a threat to someone else's autonomy. Mantovani et al. (2014) put it as follows:

In addition, it must be pointed out that apps mobile devices are not used by abstract individuals, but by people with flesh and bones, different levels of understanding and even different capacities for the exercise of individual autonomy. The ability of an individual to be able to gauge truly the exact nature of his/her situation in an mHealth environment will vary enormously between people such as a teenager or an elderly patient. In a real-life environment (in a hospital, for example) a healthcare provider would be able to guide users/patients through the process of consent, explain the consent form that needs to be signed and to answer possible questions. Current medical apps often leave the user alone and even require him/her to open up additional links to find information on external sites" (ibid., 57).

Furthermore, as pointed out by Mittelstadt et al., "While autonomy is increased by the release of the lifeworld from the confines of hospitalisation, PHM still allows the

system to invade the lifeworld and exert control through the quantisation and regulation of behaviour in the personal environment. ... The visibility of the PHM may also affect the user's identity, as PHM use becomes part of who they are, and affect behavioural patterns derived from the lifeworld. Behavioural patterns must be adapted to meet the requirements of PHM, whether that is in routines of monitoring by recording and transmitting physiological and behaviour data, or by routine of intervention, where therapies are conducted in response to the output of the PHM" (Mittlestadt et al. 2014, 50).

Recent attempts to increase health literacy, particularly in chronic conditions, have led to a shift from patients that once were inactive sufferers to active, competent partakers in their own health care. As such, patients become "knowledgeable about their condition, health services and their rights as a patient; skilled and organised in self-managing it; actively involved in information seeking and use; communicative with health professionals in an assertive manner; able to seek and negotiate treatment options" (Edwards et al. 2012, 6). Users with such a level of health literacy that use mHealth apps will be very much in control of the way they use the app and will have a clear idea of what to expect from the app.¹⁰

Having cleared some conceptual ground as to how agency and patiency are deeply intertwined, I now turn to put into perspective the widely held conviction that persuasive mHealth devices primarily emphasize potential characteristics of people's lives. Removing the hurdle of the agential bias, my goal is to show that this potential emphasis needs not to have a negative connotation. Rather, as I argue in what follows, patiency can be an extension of the agent's own mind and volition—perhaps even widening the scope of their autonomy. In making the case for complementing agency with patiency in mHealth apps, I suggest some strategies that help preserving and ultimately widening agents' autonomy in this context.

4 Volitional Aids: Enhancing Agency by Design

MHealth apps that are aimed at persuading users to change their behavior have philosophically been challenged mainly on two related grounds:

- (a) Based on eroding users' autonomy due to their persuasive character. Rossi and Yudell, for example, claim that "persuasive (as opposed to manipulative) health communication infringes upon autonomy if and when it exerts a controlling influence, and persuasion may infringe upon autonomy if risk or health messages fail to provide message recipients with the information they are due" (Rossi and Yudell 2012, 201).
- (b) For addressing users primarily as patients. Sharon (2017) as previously mentioned, sees the potential control over users enabled by mHealth apps as a threat to their autonomy, degrading their level of agency towards inactive patients.

Sharon (2017), as previously mentioned, sees the potential control over users enabled by mHealth apps as a threat to their autonomy, degrading their level of agency towards inactive patients.

Now, at first glance, it seems evident that being persuaded by a device that tells us how to conduct significant portions of our lives renders the agent passive, an inert

¹⁰ Thanks to an anonymous referee for drawing my attention to this.

recipient of technological commands. If these intrusive commands are, for good measure, disguised as persuasive suggestions (rather than straightforward imperatives), and thus chip away at users' autonomy, employing such devices should be avoided at all costs. Or so it seems, given the agential bias.

Heretofore, I have tried to show that persuasion is neither necessarily paternalistic nor inherently at odds with autonomy and that being acted upon is an inevitable, ever present part of agents' lives—not something to be evaded, as those who hold agency dear might think. I am now in a position to look at the philosophically sometimes underappreciated positive side of mHealth apps. In so doing, I bring together the extended mind and extended will theses with a more balanced account of agency that encompasses both agential and patential characteristics, contending that there is a way of technologically harnessing patency to enhance agency.

4.1 Extended Minds and Extended Wills

One contentious claim that, perhaps, sparks the critique of mHealth apps as agency-eroding is the juxtaposition of users on the one hand and technological devices on the other. This sharp division between agents as bearers of mental states and technology has not been made explicit in the context of mHealth apps (to my knowledge) but might, perhaps, be an implicit reason why apps are seen as mere tools that are not part of the agent's cognitive or volitional apparatus. This follows both straightforwardly from traditional conceptions of the mind such as physicalism and from recent criticisms of the extended mind thesis that see the mind as staying "safely within the boundaries of the body and brain" (Weiskopf 2008, 275). mHealth apps as tools of external influence could then, if they were to impose someone else's interest on users, present a threat to their autonomy. This traditional boundary between agents and their environment has been called into question ever since Andy Clark's and David Chalmers's extended mind thesis (Clark and Chalmers 1998). On this view, certain technological devices are literally extensions of people's minds, enabling agents to extend their minds beyond the physical boundaries of their bodies. What kind of cognitive processes qualify as being realized extendedly depends on a sensible conditional:

If, as we confront some task, a part of the world functions as a process which, *were it done in the head*, we would have no hesitation in recognizing as part of the cognitive process, then that part of the world *is* (so we claim) part of the cognitive process. Cognitive processes ain't (all) in the head (ibid., 8; italics in original)!

Clark and Chalmers distinguish between what they call "extended cognition" which involves various physical and computational artifacts employed by agents, such as calculators and notebooks. These are enactive systems that are not confined to the physical boundaries of one's body, but nevertheless an extension of the agent's cognitive apparatus into the environment. Functionally, extended cognition plays the same role as does internal cognition, or as in the calculator example, extended cognition complements, even enhances internal cognition. For example, most of us are reasonably decent at mental arithmetic but nowhere near the performance of a calculator.

A stronger and more controversial thesis is what Clark and Chalmers describe as “extended mind”—the claim that some mental states, particularly beliefs and desires, can literally be stored and manifested on devices outside of one’s own body without thereby ceasing to be one’s own beliefs and desires. Proponents of this view take smartphones to be apt examples of mind extenders. Not only do these devices complement the agent’s internal cognition, as say in their functions as calculators, but they also do store information such as pictures, call logs, and directions that constitute the agent’s own mental states. Take, for example, finding one’s way around in a familiar but not so frequently visited area. When recalling the way to get from here to there, it makes no difference, so says the extended mind thesis, whether we reach the destination by accessing our internal memory or by consulting our smartphone’s memory that has that information stored on our behalf. Either way, we are recollecting an existing belief.

Some authors have expanded the extended mind theory to an outsourcing of decision-making capacities, which they call “extended will” (Heath and Anderson 2010). On this view, people make use of the ability to offload various motivational and cognitive processes to their environment, broadly construed. Due to such outsourcing, the environment can provide the necessary “scaffolding” (Sterelny 2010) that enables agents to successfully solve various problems of self-regulation whose accomplishment by traditional, internal means might have been impossible due to a temporary or permanent scarcity of the corresponding internal resources. From the extended will perspective, persuasive technologies do not appear as a threat to agents’ autonomy but rather as a form of what I call volitional aids, assisting agents with the accomplishment of difficult tasks. Such technologies, seen in this light, do not undermine agents’ autonomy because they are genuinely parts of our own decision-making processes. Just as outsourced beliefs are genuine parts of the extended mind, so are outsourced volitions genuine parts of the extended will.

Granted that both the extended mind and the extended will theses are hotly debated, they nevertheless present a sensible platform for suggesting that mHealth apps and an agent’s internal cognitive and volitional apparatus need not be mutually exclusive opposites. Particularly, the extended mind/will theses call into question the claim that mHealth apps are agents of their own, either in virtue of imposing the app designer’s will on users or by representing some other agent’s vested interest by proxy.

One might argue that there is no tight conceptual connection between the agent/patient distinction and the extended mind/will discussion.¹¹ This might be so, but realizing that patiency is an integral part of every agent’s life—even though it is mainly constituted by things that happen to us—helps understanding why external devices such as mHealth apps that occasionally render the agent passive in a similar way by, for example, nudging users, can nevertheless be genuine parts of the agent’s cognitive or volitional apparatus. The very fact that things happen to us that are beyond our direct initiation does not necessarily render these happenings foreign.

In what follows, I take the thesis that mHealth apps are serious candidates for both mind and will extenders as a tenable way of understanding the relation between agents and these kinds of systems. This leaves me with the question of what features such

¹¹ I owe this point to an anonymous referee.

devices must have in order to initially at least retain users' autonomy and to potentially even enhance it.

I now turn to discuss three principles of persuasion that shall help rendering mHealth devices ethically sound and spell out how an agency enhancement via persuasive mHealth apps can work by way of looking at some real-life examples.

4.2 Ethically Sound Principles of Persuasion

Spahn (2012) who sees merit in persuasion suggests three useful principles that ensure the preservation of as much agency as possible while at the same time taking advantage of the effectiveness of technological persuasion that enables users to reach goals more efficiently.

- (1) Persuasion should be based on prior (real or counterfactual) consent.

Any persuasive device that even initially bears the potential of preserving users' autonomy must meet the gold standard of obtaining informed consent before the device is put to use. Once it is ensured that users apprehend and consent to the device's persuasive character, Dworkin's second criterion of paternalistic interventions (*X* does something to *Y* without their consent) is circumvented. Practically, this could work for example by presenting users' educational videos that sincerely display the device's workings and persuasive goals. If this is implemented as a prerequisite for being able to operate the device, informed consent is warranted.

- (2) Ideally, the aim of persuasion should be to end persuasion.

Persuasive technologies involve a specific kind of human-technology interaction that differs from regular human-to-human communication in at least two important ways: both parties have limited resources to influence each other, and there is no mutual communication possible in the sense that the device does not understand users' feelings and thus cannot adequately respond to them. The interaction between users and persuasive technologies can thus be characterized as an asymmetrical relation with the goal of changing users' behavior for the better. When it comes to the means of persuasion, it is important to distinguish between what might be called *manipulative* persuasion and *educational* persuasion. Manipulative persuasion aims at creating a dependent person that is in permanent need of guidance. The aim of educational persuasion, on the other hand, is to empower users and thus to promote their independency and autonomy. The autonomous user is, then, able to end the asymmetrical relation and to educate themselves. For example, if an mHealth app educates users to think about their nutrition behavior and helps implementing a healthier diet, eventually the persuasive technology is no longer needed, and users will be able to stick to their newly acquired routine by themselves.

- (3) Persuasion should grant as much autonomy as possible to the user.

Persuasive technologies preserve users' autonomy just in case that these devices do not take over users' large-scale decision-making capacity. This issue is particularly tricky since one major asset of persuasive technologies is precisely to take over some of

users' choices, prompting them to follow behavioral recommendations generated by the device. Now, the key to autonomy preservation and, ultimately, to autonomy enhancement is to ensure that large-scale goals of behavior change are set autonomously.

4.3 Volitional Aids and Second-Order Autonomy

Preserving and ultimately enhancing users' autonomy might be achieved by what I have earlier described as volitional aids. If the overall goals of behavior change are set by users themselves (i.e., autonomously), there is no other agent involved and thus no threat to autonomy. The device merely aids users' initial intentions to achieve their goals via technological persuasion. To further embellish this idea, I differentiate between first- and second-order autonomy:

First-order autonomy can, in this context, be described as an agent's exerted capacity to autonomously make decisions on a small-scale level. For example, deciding at whim what to have for lunch, how to get to work, whether to hit the gym today.

Second-order autonomy can, in this context, be described as an agent's exerted capacity to autonomously make decisions on a large-scale level. For example, deciding to improve one's diet, or to live a more active life by increasing one's sports activities.

In the spirit of Harry Frankfurt's hierarchical model of autonomy, agents are autonomous with respect to their actions if and only if their first-order autonomous small-scale decision-making capacity is approved of (or sanctioned by) their second-order autonomous large-scale decision-making capacity. Second-order autonomy oversees first-order autonomy, as it were.

The suggested view of second-order autonomy as an agent's exerted capacity to autonomously make decisions on a large-scale level that might include the occasional forfeit of first-order autonomous choices combined with the hypothesis that mHealth apps can be seen as extended minds/wills potentially sheds new light on Spahn's second principle of ethically sound persuasion as ultimately aiming to end persuasion.¹² If I am correct in conceptualizing mHealth apps as volitional aids, there is no need to aim for an end of persuasion. By using such volitional aids, we simply outsource internal willpower to the environment that nevertheless remains part of the agent's own volitional apparatus which is based on a second-order autonomous choice to employ such technologies.

It is important to keep in mind that mHealth apps are not imposed on users but autonomously employed (except, perhaps, in a few medical settings). That said, mHealth apps do not only target people with health conditions but are also frequently used by healthy people and created for prevention purposes. Lifestyle-based interventions such as motivational goal setting, action planning, or self-monitoring are becoming more feasible for individuals due to personalized mobile technologies (Orrell and Brayne 2015). Recent reviews suggest that self-monitoring applications have great potential to aid and modify people's lifestyle (Burke et al. 2015) and to encourage self-management in chronic conditions and patient autonomy (Boulos et al. 2011; Landry 2015). First encouraging effects with respect to the use of such apps have been demonstrated with respect to lifestyle issues such as physical activity, diet, and weight

¹² I am grateful to an anonymous referee for pointing me in that direction.

control (Carter et al. 2013; Glynn et al. 2014; Lubans et al. 2014). Both recreationally and medically, mHealth apps are mainly used for tracking and monitoring purposes only or additionally for helping users to change their behavior. When employed with the explicit goal of being assisted to change one's behavior, users expect such devices to persuade them towards a wanted outcome.¹³ In these circumstances, persuasion can hardly count as posing a threat to autonomy, but rather as an expression of users' second-order autonomous choice to use such devices.

Now, in some cases, acting truly autonomously might mean to voluntarily relinquish or outsource one's first-order autonomy for the purpose of enhancing one's second-order autonomy. Let us return to the initial lunch example that can now be redescribed in the following way:

If I have autonomously decided that I want to improve my diet as an exercise of second-order autonomy, I will have salad for lunch even though this might be at odds with satisfying my cravings for fries as an exercise of first-order autonomy. Persuasive mHealth apps can serve as volitional aids in such scenarios since they have the potential to increase one's second-order autonomy by incentivizing a first-order autonomous behavior that is in accordance with the previously set second-order autonomous goal. In a way, then, by increasing the level of patency regarding small-scale decisions ("I'll go with what the device tells me"), the overall level of agency is enhanced, trading patency in the fine print for agency in the heading, as it were. Since the behavior change is self-initiated, and thus based on the agent's own intentions and motivations, there is no threat to second-order autonomy but an enhancement thereof. Importantly, autonomy-enhancing persuasive apps treat users as reason-responsive agents by way of presenting reasons and motivational incentives for self-initiated behavior changes rather than simple imperatives. Weintraub and Barilan (2001) go as far as to suggest that the value of autonomy traces back to persons' right to be respected as agents who can argue, persuade, and be persuaded in matters of utmost personal significance such as decisions about medical care. These authors suggest that autonomy should and could be respected only after such an attempt of persuasion has been made.

How can this work in practice? One established technique that helps changing one's behavior by effectively translating previously set goals into action are so-called implementation intentions (Gollwitzer 1999; Roughley 2016). Such "if-then plans" are psychological constructs for establishing new routines that aim at long-term behavioral changes. The basic structure of implementation intentions looks as follows:

If situation x arises, I will initiate the goal-directed response y .

Whereby x constitutes the if-component, representing a critical situation containing behavioral cues. Y constitutes the then-component, representing the goal-directed behavior. For implementation intentions to work most effectively, the striven-for plans must be both viable and precise.

Coming back to the previous example, an implementation intention that promotes the agent's second-order autonomous goal to improve their diet could be the following conditional: "If there is an healthy option at the cafeteria for lunch, I will go for it." If

¹³ Thanks to an anonymous referee for pointing this out.

mHealth apps are used to remind users of that intention and the corresponding cue by, say, popping up a message at noon, so much the better for their effective goal achievement. Implementation intentions facilitate second-order autonomy through harnessing potential features, namely following previously set goals by hewing to persuasive suggestions. “The device tells me to do so, so I’ll comply”—thereby making the achievement of the previously set second-order autonomous goals more efficacious. Setting such large-scale goals preserves autonomy, and technological persuasion makes their achievement more effective.

It is, however, crucial to keep some caveats in mind when looking for practical solutions to make persuasive mHealth apps work. Importantly, the goals set by second-order autonomous decision making capacities must be self-motivated in order to have their intended effect. According to the “self-determination theory” (Ryan and Deci 2000), “autonomous motivation” is much more effective by way of sustainably changing one’s behavior compared to what Ryan and Deci call “controlled motivation.” When agents are autonomously motivated, they gain self-support and reinforcement through their own actions; the motivation emanates from the self, and the behavior is thus self-determined (Hager et al. 2014, 567). On the other hand, controlled motivation is an external, introjected regulation of one’s behavior (e.g., avoiding punishment or feelings of guilt). There is ample evidence suggesting that autonomous motivation has the most pervasive effects on behavior change, particularly on health-related behavior (ibid., 578). Pavey and Sparks (2010) further show that autonomous motivation increases an healthier lifestyle by promoting intentions to reduce behavior that is harmful to one’s health.

5 Concluding Remarks

In this paper, I have tried to show that persuasive mHealth applications are, despite appearances, not necessarily at odds with users’ autonomy. This is so, I have argued, for two main reasons. (1) Once the misguided assumption of a sharp divide between agency and patiency is mitigated, it becomes clear that displaying agency can be extended to potential characteristics of our lives. For example, complying with an apps’ suggestion to bike to work instead of taking the car, notwithstanding one’s current lazy preference for driving to work, might initially appear to render the agent a patient, a passive recipient of technological commands. However, harnessing this potential feature can be fully compatible with one’s autonomy if it is an exercise of adhering to a previously set large-scale autonomous goal such as wanting to increase one’s physical activity. (2) Drawing on the extended mind and extended will theories, I have argued that ethically sound persuasive technologies do not constitute intrusive external interventions into people’s lives but are rather what I have called volitional aids, assisting agents with the accomplishment of difficult tasks that might have been impossible to achieve otherwise due to a temporary or permanent scarcity of the corresponding internal resources. Ethically sound persuasive apps, thus, need not be paternalistic and can even bear the potential to enhance agents’ autonomy when applied with caution.

Acknowledgements Special thanks to the audience and the referees of the 20th Conference of the Society for Philosophy and Technology (SPT) in Darmstadt for their helpful comments on earlier versions of this paper. Thanks also to the referees of *Philosophy and Technology* for their valuable feedback.

Funding Information I am grateful to the German Federal Ministry of Education and Research (BMBF) for their generous financial support.

References

- Asimakopoulos, S., Asimakopoulos, G., & Spillers, F. (2017). Motivation and user engagement in fitness tracking: heuristics for mobile healthcare wearables. *Informatics*, 2017, 1–16.
- Boulos, M. N. K., Wheeler, S., Tavares, C., & Jones, R. (2011). How smartphones are changing the face of mobile and participatory healthcare: an overview, with example from eCAALYX. *Biomedical Engineering Online*, 10, 24.
- Burke, L. E., Ma, J., Azar, K. M. J., Bennett, G. G., Peterson, E. D., Zheng, Y., & Quinn, C. C. (2015). Current science on consumer use of mobile health for cardiovascular disease prevention: a scientific statement from the American Heart Association. *Circulation*, 132(12), 1157–1213.
- Carter, M. C., Burley, V. J., Nykjaer, C., & Cade, J. E. (2013). Adherence to a smartphone application for weight loss compared to website and paper diary: pilot randomized controlled trial. *Journal of Medical Internet Research*, 15(4), e32.
- Chatterjee, S., & Price, A. (2009). Healthy living with persuasive technologies: framework, issues and challenges. *Journal of the American Medical Informatics Association*, 16(2), 171–178.
- Cholbi, M. (2017). Paternalism and our rational powers. *Mind*, 126(501), 123–153.
- Cialdini, R., et al. (2005). Persuasion and health: creating positive behaviour change. In J. Kerr, R. Weitkunat, & M. Moretti (Eds.), *ABC of behavior change: a guide to successful disease prevention and health promotion* (pp. 247–258). Edinburgh: Elsevier Science.
- Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.
- Davis, R., Campbell, R., Hildon, Z., Hobbs, L., & Michie, S. (2015). Theories of behaviour and behaviour change across the social and behavioural sciences: a scoping review. *Health Psychology Review*, 9(3), 323–344.
- Dworkin, G. (1972). Paternalism. *The Monist*, 56, 64–84.
- Dworkin, G. (2005). Moral paternalism. *Law and Philosophy*, 24(3), 305–319.
- Dworkin, G. (2015). Defining paternalism. In Thomas Schramme (ed.), *New perspectives on paternalism and health care*, Springer, 17–29.
- Dworkin, Gerald, “Paternalism”, The Stanford encyclopedia of philosophy (Spring 2017 Edition), Edward N. Zalta (ed.), forthcoming URL = <<https://plato.stanford.edu/archives/spr2017/entries/paternalism/>>.
- Edwards, M., Wood, F., Davies, M., & Edwards, A. (2012). The development of health literacy in patients with a long-term health condition: the health literacy pathway model. *BMC Public Health*, 12, 1–15.
- Enoch, D. (2016). What’s wrong with paternalism: autonomy, belief and action. *Proceedings of the Aristotelian Society*, 116(1), 21–48.
- Fogg, B. J. (2003). *Persuasive technology: using computers to change what we think and do*. The Morgan Kaufmann series in interactive technologies. Amsterdam, Boston: Morgan Kaufmann Publishers.
- Glynn, L. G., Hayes, P. S., Casey, M., Glynn, F., Alvarez-Iglesias, A., Newell, J., & Murphy, A. W. (2014). Effectiveness of a smartphone application to promote physical activity in primary care: the SMART MOVE randomised controlled trial. *The British Journal of General Practice*, 64(624), e384–e391.
- Gollwitzer, P. (1999). Implementation intentions: strong effects of simple plans. *American Psychologist*, 54(7), 493–503.
- Groll, D. (2014). Medical paternalism—part 2. *Philosophy Compass*, 9(3), 194–203.
- Heath, J., & Anderson, J. (2010). Procrastination and the extended will. In M. White & C. Andreou (Eds.), *The thief of time* (pp. 233–252). New York: Oxford University Press.
- Karlsson, M. (2002). Agency and patency: back to nature? *Philosophical Explorations*, 5(1), 59–81.
- Koelle, M., Kranz, M., & Lindemann, P. (2014). Persuasive technologies and applications. *Advances in Embedded Interactive Systems*, 3(2).
- Krakauer, E. (1998). Prescriptions: autonomy, humanism and the purpose of health technology. *Theoretical Medicine and Bioethics*, 19, 525–545.
- Krieger, W. (2013). Medical apps: public and academic perspectives. *Perspectives in Biology and Medicine*, 56(2), 259–273.

- Landry, K. (2015). Using eHealth to improve health literacy among the patient population. *Creative Nursing*, 21(1), 53–57.
- Lanzing, M. (2016). The transparent self. *Ethics and Information Technology*, 18(1), 9–16.
- LeBar, M. (2013). *The value of living well*. Oxford: Oxford University Press.
- Lott, M. (2016). Agency, patiency, and the good life: the passivities objection to eudaimonism. *Ethical Theory and Moral Practice*, 19(3), 773–786.
- Lubans, D. R., Smith, J. J., Skinner, G., & Morgan, P. J. (2014). Development and implementation of a smartphone application to promote physical activity and reduce screen-time in adolescent boys. *Frontiers in Public Health*, 2, 42.
- Mantovani, E., et al. (2014). eHealth to mHealth—a journey precariously dependent upon apps? *European Journal of ePractice*, 20, 48–66.
- Mittlestadt, B., Fairweather, B., Shaw, M., & McBride, N. (2014). The ethical implications of personal healthcare monitoring. *International Journal of Technoethics*, 5(2), 37–60.
- Nordgren, A. (2015). Privacy by design in personal health monitoring. *Health Care Analysis*, 23, 148–164.
- O’Keefe, D. J. (2012). Conviction, persuasion, and argumentation: untangling the ends and means of influence. *Argumentation*, 26(1), 19–32.
- Orrell, M., & Brayne, C. (2015). Dementia prevention: call to action. *The Lancet*, 386(10004), 1625.
- Owens, J. & Cribb, Alan (forthcoming). ‘My Fitbit thinks I can do better!’ Do health promoting wearable technologies support personal autonomy? *Philosophy and Technology*: 1–16.
- Pavey, L. J., & Sparks, P. (2010). Autonomy and reactions to health-risk information. *Psychology and Health*, 25(7), 855–872.
- Reader, S. (2007). The other side of agency. *Philosophy*, 82(4), 579–604.
- Rossi, J., & Yudell, M. (2012). The use of persuasion in public health communication: an ethical critique. *Public Health Ethics*, 5(2), 192–205.
- Roughley, N. (2016). *Wanting and intending: elements of a philosophy of practical mind*. Dordrecht: Springer.
- Sharon, T. (2017). Self-tracking for health and the quantified self: re-articulating autonomy, solitary and authenticity in an age of personalized healthcare. *Philosophy & Technology*, 30(1), 93–121.
- Spahn, A. (2012). And lead us (not) into persuasion...? Persuasive technology and the ethics of communication. *Science and Engineering Ethics*, 18(4), 633–650.
- Sterelny, K. (2010). Minds: extended or scaffolded? *Phenomenology and the Cognitive Sciences*, 9(4), 465–481.
- Taylor, R. (1982). Agent and patient: is there a distinction? *Erkenntnis*, 18(2), 223–232.
- Timmer J., Kool L., van Est R. (2015) Ethical challenges in emerging applications of persuasive technology. In: MacTavish T., Basapur S. (eds) Persuasive technology. PERSUASIVE 2015. Lecture Notes in Computer Science, vol 9072. Springer, 196–201.
- Weintraub, M., & Barilan, M. (2001). Persuasion as respect for persons: an alternative view of autonomy and of the limits of discourse. *Journal of Medicine and Philosophy*, 26(1), 13–34.
- Weiskopf, D. A. (2008). Patrolling the mind’s boundaries. *Erkenntnis*, 68(2), 265–276.