

Digital Wellness and Persuasive Technologies

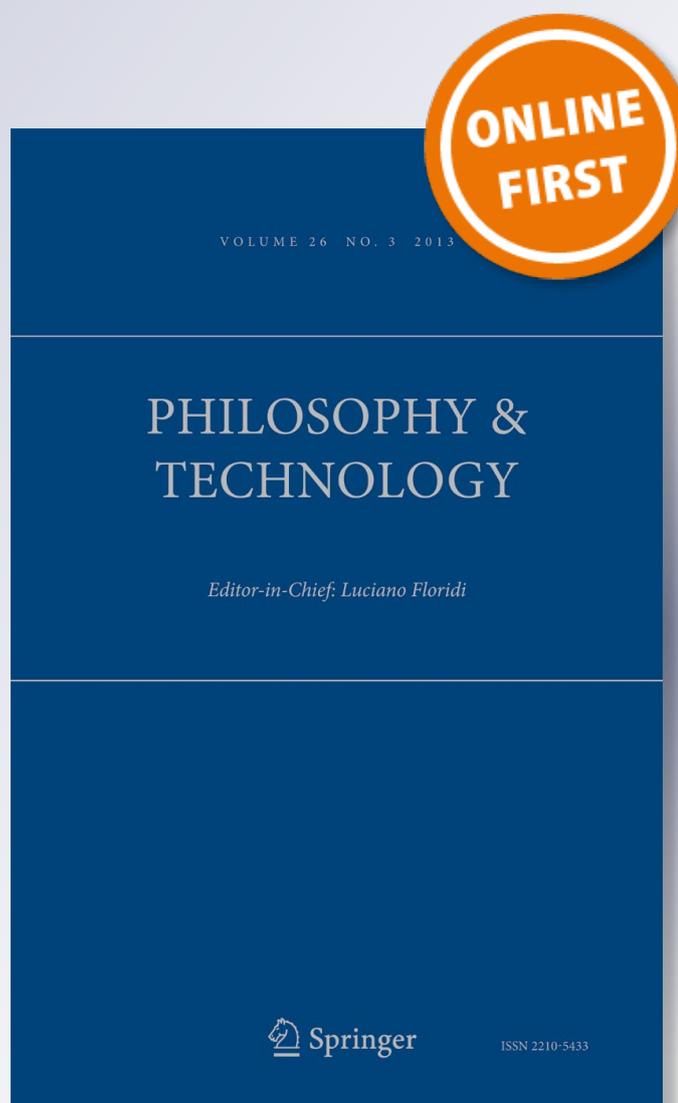
Laura Specker Sullivan & Peter Reiner

Philosophy & Technology

ISSN 2210-5433

Philos. Technol.

DOI 10.1007/s13347-019-00376-5



Your article is protected by copyright and all rights are held exclusively by Springer Nature B.V.. This e-offprint is for personal use only and shall not be self-archived in electronic repositories. If you wish to self-archive your article, please use the accepted manuscript version for posting on your own website. You may further deposit the accepted manuscript version in any repository, provided it is only made publicly available 12 months after official publication or later and provided acknowledgement is given to the original source of publication and a link is inserted to the published article on Springer's website. The link must be accompanied by the following text: "The final publication is available at link.springer.com".



Digital Wellness and Persuasive Technologies

Laura Specker Sullivan¹ · Peter Reiner²

Received: 26 December 2018 / Accepted: 1 September 2019 / Published online: 07 September 2019
© Springer Nature B.V. 2019

Abstract

The development of personal technologies has recently shifted from devices that seek to capture user attention to those that aim to improve user well-being. Digital wellness technologies use the same attractive qualities of other persuasive apps to motivate users towards behaviors that are personally and socially valuable, such as exercise, wealth-management, and meaningful communication. While these aims are certainly an improvement over the market-driven motivations of earlier technologies, they retain their predecessors' focus on influencing user behavior as a primary metric of success. Digital wellness technologies are still persuasive technologies, and they do not evade concerns over whether their influence on users is ethically justified. In this paper, we describe several ethical frameworks with which to assess the justification of digital wellness technologies' influence on users. We propose that while some technologies help users to complete tasks and satisfy immediate preferences, other technologies encourage users to reflect on the values underlying their habits and teach them to evaluate their lives' competing demands. While the former approach to digital wellness technology is not unethical, we propose that the latter approach is more likely to lead to skillful user engagement with technology.

Keywords Persuasive technology · Digital wellness · Paternalism · Maternalism

1 Introduction

In the early days of smartphone application development, developers focused on capturing user attention. The more clicks and likes and views, the better, creating what has come to be called the attention economy, driven by an ecosystem of apps that range from games to social media to email. These so-called persuasive technologies were designed to “[develop] new digital experiences that influence people” (Fogg 2009;

✉ Laura Specker Sullivan
speckersullivanle@cofc.edu

¹ College of Charleston, 66 George Street, Charleston, SC 29424, USA

² University of British Columbia, 2329 West Mall, Vancouver, BC V6T 1Z4, Canada

Larson 2014; Byrnes 2015). Persuasion as used here is a generic term, referring simply to the fact that these technologies, when successful, influence the behavior of users.¹ Behavior change can occur through multiple avenues: through rational and non-rational means that include coercion, manipulation, addiction, and rational argumentation. As other commentators have noted, whether persuasive technology is ethically sanctioned depends not on the means of influence itself but on the circumstances in which the technology will be used and the reason for its use (Spahn 2012).

As society has come to terms with the cost of this influence—distracted driving, unfocused conversations, scarce opportunities for contemplation, among others—application developers have embraced digital wellness as a product. Digital wellness technologies often use strategies similar to those used in persuasive apps to motivate users towards behaviors that are personally and socially valuable, such as exercise, wealth-management, and meaningful communication. Some apps even focus exclusively on helping people to use apps and personal technologies with less frequency.

Recently, ethicists have expressed concern that certain persuasive technologies—even those aimed at health or wellbeing—may be ethically suspect in that they limit, threaten, or compromise autonomy (Verbeek 2009; Nagel et al. 2016; Owens and Cribb 2017; Wagner 2018; Lanzing 2018; Burr et al. 2018). The worry is that these apps supersede individuals' liberty to make lifestyle choices for themselves and that they damage individuals' autonomous capacity to make these choices well. Yet autonomy is not the final word in the ethical justification of technological interventions; wellbeing is valued as well, and on some accounts, the two cannot be so easily pulled apart. In the context of digital wellness, it is not always clear which among the range of possible human goods technologies are aiming to achieve with their interventions.

In this paper, we systematically examine some of the different interventions used by digital wellness technologies and highlight their likely goals. We use this examination to analyze ethically both the impact that individual apps can have on users' decision-making and the broader societal context within which these decisions are made. We begin by describing the different philosophical frameworks within which to understand the types of influences that persuasive technologies represent. We then use a number of exemplar technologies to categorize persuasive technologies according to these conceptual frameworks. Finally, we use the pairing of our exemplar technologies and conceptual frameworks to analyze the ethical dimensions of the persuasion employed by digital wellness technologies.

2 Frameworks of Influence

There exists a range of frameworks with which to conceptualize the ethics of influence on individuals' decisions and actions. The most expansive framework is paternalism, which describes interventions in which an individual's liberty or autonomy is compromised for the sake of their wellbeing. There are multiple ways to define paternalism and many different types of paternalistic action. In a governmental context, paternalism

¹ Persuasive technologies may be conceptually distinguished from technologies aimed at behavior change (Smids 2018). In this article, we use the two terms interchangeably in order to address both types of technological influence.

might take the form of a public policy, such as requiring those riding in a motor vehicle to wear a seatbelt. In an interpersonal context, paternalism might be a decision one person makes for another. Paternalism can even be architectural, as when an office building is designed to encourage people to take the stairs rather than the elevator, for example by placing the stairs closer to the door. Similarly, in the context of technology, paternalism is often described in terms of the design of choices permitted by the user. What these different forms of paternalism have in common is that they indirectly structure or directly interfere with people's decisions and behavior based on the paternalist's conception of what is in the best interests of the user. Yet within these diverse paternalisms, conceptual frameworks differ most starkly along the lines of the influence, the motivation of the influence, and the goals of the influence, as we describe below.

2.1 Classic Paternalism

Classic or “hard” paternalism is so-called due to the heavy-handed nature of the intervention. A common example is when a physician withholds a diagnosis from a patient due to the suspicion that the truth of the diagnosis will harm them in some way. Many governmental policies fit this description as well, as with laws requiring seat belts, restricting the size of soda bottles, and so on. These paternalistic decisions not only do not take individuals' wishes into account, they may even be made in direct opposition to their expressly stated wishes (Groll 2012; Dworkin 1972).

The classic conceptual formulation of hard paternalism is that it occurs when an individual interferes with another individual without their consent and for their benefit (Dworkin 1972). There have been many reformulations of this concept, with one of the more influential being that paternalism is characterized best by the paternalizing agent's mistrust of the object of paternalism's ability to successfully make beneficial decisions for him or herself. According to this form of paternalism, what matters more than the paternalist's actual violation of the object's liberty is the paternalist's motive in acting paternalistically—replacing the object's judgment with his or her own (Shiffrin 2000; Dworking 2013; Begon 2016). What is captured by this mistrust is the infantilizing nature of paternalism—when an individual's ability to make their own decisions about their welfare is mistrusted and then superseded, this implies that they cannot make these decisions for themselves. This replacement of the individual's decision-making not only violates liberty immediately but may even have downstream effects on the individual's own conception of their autonomous capacities, as they may begin then to mistrust their own decision-making abilities (Conly 2013; Christman 2014).

2.2 Libertarian Paternalism, or Nudging

An amendment to such hard paternalism has been offered by Thaler and Sunstein in the form of libertarian paternalism (Thaler and Sunstein 2009). In this formulation, the paternalist is renamed as a choice architect who structures the environment in such a way that individuals are more likely to make the choice that the libertarian paternalist judges to be best for them, yet individuals are still free to choose any outcome they prefer (Thaler et al. 2010). Such libertarian paternalist “nudges” have become commonplace in Western society in the last decade.

Libertarian paternalism arose from the realization that humans are subject to a number of cognitive biases—common errors of reasoning—that often lead to suboptimal decision-making (Kahneman 2011). In order to offset these errors, the libertarian paternalist alters the environment in some way so that the error is less likely. This has recently been described as “curated motivation” (Jennings et al. 2016), a term that succinctly encapsulates the essence of a nudge.

Curated motivation is distinct from traditional economic theories, which considers humans as rational agents whose behavior is primarily influenced through the manipulation of incentives such as costs and benefits. An example of persuasion on the traditional model would be a cost attached to using additional towels during a hotel stay. Visitors would be less likely to use additional towels if they have to pay more to do so. An example of curated motivation would be a notice in a hotel bathroom that tells patrons that 60% of visitors re-use their towels, thereby relying upon the subtle subconscious social desirability bias to motivate behavior rather than cost. Both models for persuasion operate at the level of motivation and desire—but they differ in their conceptions of what is actually motivationally effective. For traditional economics, cost-based incentives shape motivation. For behavioral economics, features that contribute to subconscious cognitive processing affect motivation.

In contrast to the classic paternalism described above, libertarian paternalism is meant to describe ways in which (primarily) institutions can structure decision-making to benefit their members without violating their liberty—thus the use of the term “libertarian” (Sunstein and Thaler 2003). As with the potentially infantilizing downstream results of classic paternalism, libertarian paternalism has more global effects on individuals as well. As Mitchell points out, one of the key premises of libertarian paternalism is that individuals will act irrationally in making certain choices, and thus they are not always the best judges of their own welfare (Mitchell 2005). Thus, libertarian paternalism has essentially the same premise as classic paternalism—a mistrust of individuals’ abilities to make good choices for themselves. Rather than work to build up these individuals’ autonomy competencies with regard to these particular preferences (where autonomy competencies include the abilities to critically reflect on past decisions, forecast the effects of future choices, and reflect on motives and desires (Meyers 1989)), libertarian paternalism, like classic paternalism, seeks to make the decision that the paternalist thinks is locally better for the object of paternalism more likely without considering the effect on the individuals’ overall reflective capacities. In short, libertarian paternalism curtails individuals’ opportunities to exercise reflective decision-making practices by manipulating which choices are available and how they will appear. Of course, libertarian paternalism is not necessarily unethical, and it is meant to move people towards the outcomes they themselves would choose for themselves, if they had all the information, were effective reasoners, and were not subject to the distortive effects of cognitive biases.

2.3 Means Paternalism and Personalized Paternalism

Like libertarian paternalism, means paternalism and personalized paternalism attempt to soften the sting of classic paternalism by retaining some respect for individuals’ liberty. Libertarian paternalism offers individuals the freedom to make ultimate decisions derived from a curated set of choices in which some options are more heavily

weighted than others. Means paternalism provides individuals with opportunities to set goals and then restricts the possible means that can be used to reach those goals, while personalized paternalism seeks to fit both the means of paternalism and the goals to individuals' own preferences.

Means paternalism is exemplified by situations in which individuals set goals for themselves but require another individual's expertise to best reach those goals. For example, hiring a coach to train for a triathlon, or a nutritionist to lose weight, or a wedding planner to plan a wedding are all cases of means paternalism. Individuals in each case know what they want—to podium in triathlon, or to lose a certain amount of weight, or to have a seamless wedding, but they need others to tell them what to do in order to get there. In these cases, the paternalistic interventions are extremely circumscribed, or local: the concern is with reaching individual preferences or goals, and are judged to be successful based on whether or not the agent meets those goals. In other words, autonomy is decision-relative. Yet at the global level—that is, the broader picture of the individual's overall autonomy competencies—there is not necessarily an effect on the individual's capacities, either in terms of an infantilizing suspicion of the individual's abilities or an empowering support for the individual's reflective capacities.

Personalized paternalism proposes that the paternalist may be best at acting in the object of paternalism's best interests when the means of interference as well as the goals are suited to the individual's personal preferences and goals. Yet it is not clear why personalized paternalism remains a form of paternalism, as arguably the conceptual core of paternalism is that it is an intervention meant to benefit the person who is the object of the intervention, yet it is an intervention that restricts, contravenes, or disregards the liberty or autonomy of this person. Personalized paternalism maintains the first criterion of paternalism—acting for someone's benefit—but in discarding the second, it also loses the second criterion for paternalism. As we point out below, personalized paternalism looks more like a different conceptual framework—maternalism.

2.4 Maternalism

Specker Sullivan and Niker have argued that maternalism is a useful concept for capturing interventions that occupy the same conceptual space as personalized paternalism: “If paternalism is acting in another person's best interests without due consideration of their autonomy, maternalism is acting for the benefit of another person in a way that takes that person's autonomous agency into account, despite no explicit expression of consent or assent being given by the person on whose behalf the decision is made” (Specker Sullivan and Niker 2018; 7). The core difference is that while paternalism involves a local constraint of an individual's liberty for the sake of their immediate benefit but not necessarily their global autonomy competencies, maternalism involves a local intervention on an agent that benefits their overall autonomy competencies in conjunction with their wellbeing. Paradigmatic examples of maternalistic relationships include nurses, teachers, and coaches—individuals who care both about the wellbeing and the autonomy competencies of the individuals whom they influence.

While most forms of paternalism are conceptualized at the local level, despite their significant downstream effects on individuals' global autonomy competencies,

maternalism operates mainly on the global level. This is to say that there is no maternalistic action that can be justified based solely on a description of the particular intervention—justification of maternalistic interventions requires description of the caring nature of the relationship in which the intervention takes place as well as the history of the relationship and the basis for epistemic access of the maternalizing party. A maternalistic action, to be successful, must be in an individual's overall benefit meaning that it contributes to their welfare *and* enhances their autonomy competencies. For this reason, maternalistic interventions are inherently risky and difficult to justify, as they depend on longstanding relationships of care.

With paternalism, ethical justifiability of the intervention depends on whether or not the intervention will actually benefit the object of influence; with maternalism, ethical justifiability depends on relational and epistemic conditions (Mackenzie and Stoljar 2000). Relationally, a maternalistic intervention may be justified if the relationship is one of mutual trust between parties—this allows the influencer to act on behalf of the influencee without explicit consent. Epistemically, the relationship must be one of sufficient intimacy and duration for the influencer to know that the intervention will both benefit the influencee and support her autonomy.

The goal of describing maternalism conceptually is to highlight the ethically relevant features of interventions that resemble paternalistic interventions, yet have different justification conditions.² More detailed defense of maternalism as a useful ethical concept can be found elsewhere (Specker Sullivan and Niker 2018); here, we are primarily concerned with using maternalism as a tool with which to analyze changes in the means of persuasion that digital wellness technologies employ. In the following section, we describe how some digital wellness apps seem to be better categorized as maternalistic than paternalistic, and we explain the effect this has on an analysis of their ethical justification.

3 Categorizing Technological Persuasion

To explore the distinction between different types of persuasive technologies, consider the iPhone's settings for notifications. All apps create push notifications (in that they are "pushed out" to the user) that often both pop up on the phone's screen and cause it to vibrate, ensuring that the notification attracts the user's attention. In most cases, users accept these notifications when downloading an app, although such settings can be easily changed. Such choice architecture is explicitly designed to influence the user's behavior, drawing attention away from the task at hand and directing it to some event for which the app is an intermediary. Usually, these notifications do not even meet the requirements for paternalism, as they do not have user benefit in mind. Rather, their goal is to increase the amount of time the user pays attention to the app, no matter the cost such attention might have for the user. This was the original goal of persuasive technology—to increase both gross numbers of users and the time they spend using an app (Fogg 2009).

² While maternalism carries gendered connotations of mothering as opposed to fathering, we do not intend to import essential assumptions about gender into our analysis. Rather, just as "paternalism" as a concept has become decoupled from gender in the philosophical analysis, so we aim to take the same approach with maternalism.

With the advent of digital wellbeing, apps have been turning away from such explicitly revenue-generating persuasion and in the direction of user influence that provides more tangible benefit. Few of these technologies operate in the heavy-handed manner of hard paternalism—they do not make decisions directly for users without allowing some possibility for users to manipulate settings.

An example of libertarian paternalism in persuasive technology aimed at digital wellbeing is Gmail's aptly named new "nudge" feature (2018), which moves recently received emails to the top of the user's inbox with a reminder in red that it was received three days ago, and a question asking whether the user would like to respond. According to Gmail: "emails you might have forgotten to respond to will appear at the top of your inbox." The nudge feature is now a default feature of Gmail, and there was no opportunity for users to approve of its addition to Gmail's functionality, although it can be turned off through Gmail's settings—if a user knows how to find them. Gmail's nudge feature does not change the choices available. Users can always reply, or decide not to reply, to emails in their inboxes. It just makes it more likely that users will reply to an email that it brings to the top of the email list.

Apple's 2018 iOS update which allows users to set "downtime," or time away from their phone, is a good example of means paternalism. The goal of downtime is for the phone user to use certain apps and their phone less often, and not to use it at all during certain pre-set periods. By selecting their preferences for app and phone usage, the user indicates their goal—to use their phone less—and the phone's settings restrict their access to the phone during that time. Changing the downtime settings on the iPhone is simple, as easy as turning Bluetooth or Wi-Fi on or off. The operating system thus helps the user to reach their goals through technological means on the phone itself.

Finally, consider the app Moment (Moment.io) which is meant to help users gain control over the time they spend on their phones. It tracks phone usage (number of pickups and overall screen time) through settings already present on the iPhone. It then sets a threshold for these parameters based on past behavior, and sends alerts notifying users of their pickups and screentime for a given day so far. The goal is for users to be influenced by Moment's alerts towards using their phones less, so it operates as a form of feedback—a kind of smartphone proprioception—letting users know when to curb their usage.

Moment resembles personalized paternalism and maternalism—it uses information about user choices to encourage behavior that is most likely better for them, and also most likely what users want, given that they downloaded the app in the first place. What is interesting about Moment is that it does not supersede users' autonomous choices, but rather reminds them of their choices and gives them an opportunity to consider whether these choices are, all things considered, good for them. For this reason, Moment seems to rely less on the conceptual framework of paternalism and more on maternalism, given that the goal is to encourage users to pay more attention to the things that they signal that they care about just by downloading the app, paired with an understanding of the user's behavior (how and when they use their phone, what apps they have been spending time upon, how it compares with previous usage, etc.). In doing so, Moment does not make choices for the user nor does it manipulate the space in which the user makes choices. Rather, it stimulates user reflection based on an understanding of users' particular behavior.

The basic version of Moment is free, but one can also subscribe, for a cost, to a coaching function which offers users courses that help them attain better sleep, an

increased attention span, a family plan for screen use, and a phone use bootcamp. The last course, for example, sets different tasks for each day that the user must complete in order to advance to the next day, such as sleeping with the phone in a different room, or not bringing the phone to the bathroom. The goal of these coaching tools is not to create a desire to use one's phone less, but to help users develop the habits they need to resist the phone's temptation. Both the free version of the Moment app and the subscription version of the coaching courses give users the impression that Moment's developers care about the lives of its users. Furthermore, the transparency with which Moment works—both in free and subscription modes—and the alignment of user objectives with the app's functionality may serve to increase users trust in the app's influence, the two key criteria for maternalism.

Contrast Moment with the digital wellness app Goals,³ a newly added functionality of Google's calendar app (many of these features arose from Google's acquisition of the app Timeful, developed by behavioral economist Dan Ariely⁴). Goals will find time in users' schedules for activities that are valuable to them and insert them into their calendars. For example, a user might indicate that they want to meditate everyday for 10 min. Goals would then insert 10-min meditation periods into free slots on the user's calendar. Like Moment, Goals bills itself as an app dedicated to helping users "make more time for what really matters." Goals and Moment are also similar in that use of the app tells the developer something about the user—that they are someone who would like to use their phone less, or to use their time throughout the day more efficiently. Yet where Moment stimulates user reflection about their phone use through feedback, Goals does not stimulate user reflection—it simply does the scheduling for the user, without, for example, creating opportunities for the user to consider why it is so difficult for them to make time for activities that they ostensibly find valuable in the first place. For this reason, Goals seems to fit best as a form of means paternalism, whereby the use of the app identifies the user's goal, and the app then utilizes different means to help the user reach that goal. But the app only acts locally, by satisfying users' preferences. With Goals, the objective is pure benefit—to find time for users to do things that matter to them; the app does not in itself encourage autonomous reflection about these choices. So, even though the premise of Goals is that users will use the feature to find time for things that easily are squeezed out by busy schedules, such as exercise or meditation, it does not encourage users to consider why some activities that they find valuable are nevertheless difficult to find the time to do. Of course, this is not necessarily problematic—for users whose goal is to maximize the efficiency in an already busy schedule, Goals is a helpful technological aid.

Moment effectively takes a more global approach to users' autonomy competencies by building reflective habits in users' everyday lives. Whereas Google Calendar's Goals and Gmail's nudge features do nothing to curb individuals' reliance on technology for organization and efficiency—if anything they exacerbate it—Moment is unique in that it provides personalized feedback that can be used towards building capacities that ostensibly make use of the app unnecessary (whether or not it actually succeeds remains an empirical question). This is the point of the conceptual framework of maternalism—that it supports individuals' development of autonomy competencies

³ <https://support.google.com/calendar/answer/6334090?co=GENIE.Platform%3DiOS&hl=en>

⁴ <https://www.fastcompany.com/3045890/google-timeful>

and their welfare on a global level, rather than fulfilling individuals' immediate goals or providing them with direct benefit. In the next section, we highlight these features in our analysis of the ethical justification of the different approaches to user influence taken by these different persuasive technologies.

4 Ethically Evaluating Technological Persuasion

Persuasive technologies need not be *prima facie* objectionable. A simple example is the timing of a bank card being released from an ATM machine. If the bank card is released after the cash and receipt, it is more likely to be left behind. For this reason, many ATMs release the bank card before dispensing cash, thereby motivating patrons not to leave the card behind. Even though this example involves a choice architect designing ATM machines so as to make certain user behavior more likely, few object to this. Indeed, a major part of Sunstein and Thaler's defense of libertarian paternalistic interventions is that choices such as the release of a bank card must be arranged somehow—why not arrange them in ways that will benefit the user? While ATM's card release timing or hotels' notifications about towel use are not seen as restrictive, other uses tip the scales from justifiable benefit to liberty infringement. Facebook's newsfeed is a great example. Facebook analyzes user behavior to determine which stories users are most likely to click on, ostensibly because this is what users are most interested in seeing. It is difficult to change the algorithm that determines what is shown other than by "unfollowing" other users. Many users have become accustomed to seeing viral posts and clickbait, with few avenues for changing this display (Zuckerman 2018).

As we mention in our description of ethical frameworks above, paternalistic interventions tend to rely on a tipping of the scales between user benefit and liberty infringement. Libertarian paternalists argue that there is no liberty infringement in the case of choice architectures because choices are still available. Likewise, means paternalism and personalized paternalism describe situations in which the paternalist is merely helping other individuals to reach their goals. Seen through the lens of the range of paternalisms, many of the digital wellness technologies described above—Gmail's nudge, Apple's downtime, and Google Calendar's Goals—are not ethically objectionable, because they either encourage healthy behavior or help users reach their goals without restricting users' liberty to make choices for themselves while using the app.

This ethical assessment changes somewhat when these digital wellness apps are viewed through the lens of maternalism and contrasted with a different app, such as Moment. Nudge, Downtime, and Goals all help users to be more efficient with email correspondence, minimize smartphone use, and fit everything that matters to them into already packed schedules. What these apps fail to do is to facilitate users' reflection about the habits of responsiveness, productivity, and busyness that are quickly becoming engrained in professional life. Thus, while these apps do not violate autonomy on a local level, they do not seem to support it on a global level. That is, that they work to satisfy user preferences without encouraging them to consider why those are their preferences in the first place.

Such a criticism may seem unfair—after all, when has a technological development encouraged calm reflection on the goals and values behind its emergence? Yet given that these apps present themselves as tools of wellness and wellbeing, this is an

appropriate question to ask. After all, wellness ought to include individuals' capacities to reflect on the things that matter to them in their life and to choose their goals circumspectly and wisely, not just to keep up with the Jones's productivity levels or to squeeze in that yoga class that cannot be missed. In some ways, this is a repackaging of Susan Wolf's argument that the good life consists not just in pleasure and the lack of pain (i.e., hedonism), or getting everything that one wants, but in the construction of meaning through active engagement in projects of worth (Wolf 1997). It also aligns with Evgeny Morozov's critique of "technological solutionism," whereby merely developing a technological solution without developing the person is a thin solution to the problems of modern life (Morozov 2014).

This is where the framework of maternalism and the example of Moment can be helpful. Maternalism shifts the balance of ethical justification away from benefit as compared with liberty restriction and in the direction of a more holistic appraisal of autonomy competencies seen as congruous with wellbeing. The better we are able to reflect on our past decisions, current goals, and future wishes, the greater our capacity to ensure that we are participating actively in projects that matter to us. When our lives include justified maternalistic interventions, we include other people (and potentially technological interventions) in our lives in ways that support our reflective capacities *and* our wellbeing. While it is not unethical for some apps to focus on increasing user pleasure or satisfying their desires, this is a one-dimensional interpretation of wellbeing that leaves out users' ability to reflect on what ought to matter to them most.

Of course, maternalistic technologies are not necessarily, in and of themselves, ethical. A primary ethical question is whether users ought to feel justified in trusting these technology companies to use information about themselves for the users' benefit, and not for the companies. This is not so easy in today's technological climate, in which users are all too aware of how companies like Facebook have violated user trust in the name of profit. The premise of a persuasive technology like Moment is that users spend too much time on their smartphones, and that technology itself can be a tool to encourage users to spend less time and direct less attention to personal technology, or at least to engage with technology more skillfully.

Yet granting Moment access to a phone's internal monitoring with the apparent goal of using the phone less requires a hefty level of trust in the app developers. Users must trust that app companies care enough about users to use their data (only) for the stated goal of a healthy relationship with technology. In other words, Moment has a significant degree of access into users' habits and could easily decide to use that information for other purposes, such as advertising or market research (which there are strong financial incentives for them to do). So, users must either decide that they trust Moment's stated aims enough to hand over that access, or that the benefits of having Moment on their phone outweigh the potential costs of their data being sold to other parties (that is, assuming that the concern occurs to them at all—Draper and Turov 2019). This is arguably in stark contrast with interpersonal maternalism, where there is little reason to mistrust the larger market-based motives of individuals with whom we have longstanding and caring relationships. With new apps like Moment, there is nothing like the felt history of a teaching or a coaching relationship to fall back on in establishing trust (Andras et al. 2018). If anything, there are strong market forces and incentives that call into question any trust users may place in their apps to have their best interests, which includes their reflective capacities, in mind. This may not

necessarily be the case for apps developed by companies such as Apple that have worked assiduously to earn user trust (Tripathi 2018).

The threshold for justified trust in an app is high, and probably considerably higher than the threshold for trust in other people. But it does not necessarily mean that establishing a trusting relationship with an app is impossible—just that if apps are expected to be able to customize their influence on users and to do so in such a way that users' reflective capacities and wellbeing are promoted, users ought to be extremely cautious and app developers ought to be excessively transparent about their own motivations and aims if they would like to influence user behavior in maternalistic, and not just paternalistic, ways.

5 Conclusion

Recent debates on ethics and persuasive technology have focused on the permissibility of local influences on users' decision-making. In this paper, we have argued that the new range of technologies targeting wellbeing specifically raise the question of what this wellbeing consists in and how it is established. We have described a number of apps as focused on local or immediate effects on wellbeing without consideration of the broader effects of such interventions on users' global reflective capacities. While these interventions are not unethical, they take a narrow view of wellness that consists solely in user enjoyment or in the satisfaction of users' desires, and not in the development of users' abilities to reflect on which activities matter most to them in their lives and how they can better use technology within the context of these worthwhile activities. We have described how some apps, such as Moment, take a more maternalistic approach to users' wellbeing and concentrate on both immediate wellness and long-term reflection about wellbeing, and we have proposed that this approach is preferable and represents an aspirational ideal of digital wellness app development. After all, if persuasive technologies are to be tools that do not just influence human behavior in directions that are beneficial to companies but which human beings can use to achieve their own conceptions of the good life, then users' reflective capacities must be a component of digital wellbeing.

References

- Andras, P., Esterle, L., Guckert, M., Han, T. A., Lewis, P. R., Milanovic, K., et al. (2018). Trusting intelligent machines: deepening trust within socio-technical systems. *IEEE Technology and Society Magazine*, *IEEE*, *37*, 76–83.
- Begon, J. (2016). *Paternalism*. *Analysis*, *76*(3), 355–373.
- Burr, C., Cristianini, N., & Ladyman, J. (2018). An analysis of the interaction between intelligent software agents and human users. *Minds and Machines*. <https://doi.org/10.1007/s11023-018-9479-0>.
- Byrnes, N. (2015). Technology and persuasion. *MIT Technology Review March*, *23*, 2015 Available at: <https://www.technologyreview.com/s/535826/technology-and-persuasion/>.
- Christman, J. (2014). Relational autonomy and the social dynamics of paternalism. *Ethical Theory and Moral Practice*, *17*, 369–382
- Conly, S. (2013). *Against autonomy: Justifying coercive paternalism*. Cambridge: Cambridge University Press.

- Draper, N. A., & Turov, J. (2019). The corporate cultivation of digital resignation. *New Media & Society*, 51, 1–16.
- Dworkin, G. (1972). Paternalism. *Monist*, 56, 64–84.
- Dworking, G. (2013). Defining paternalism. In C. Coons and Weber, M. (Eds.), *Paternalism: Theory and practice*. Cambridge: Cambridge University Press.
- Fogg, B. J. (2009). Creating persuasive technologies: an eight-step design process. *Persuasive '09*. April 26–29, Claremont, California, USA.
- Groll, D. (2012). Paternalism, respect, and the will. *Ethics*, 122, 692–720.
- Jennings, B., Wertz, F., & Morrissey, M. B. (2016). Nudging for health and the predicament of agency: the relational ecology of autonomy and care. *Journal of Theoretical and Philosophical Psychology*, 36(2), 81–99.
- Kahneman, D. (2011). *Thinking, fast and slow*. Farrar, Strauss, and Giroux.
- Lanzing, M. (2018). “Strongly Recommended” revisiting decisional privacy to judge hypernudging in self-tracking technologies. *Philosophy & Technology*. <https://doi.org/10.1007/s13347-018-0316-4>.
- Larson, J. (2014). The invisible, manipulative power of persuasive technology. *Pacific Standard May*, 14, 2014 Available at: <https://psmag.com/the-invisible-manipulative-power-of-persuasive-technology-df61a9883cc7#.tg29gpms4>.
- Mackenzie, C., & Stoljar, N. (Eds.). (2000). *Relational autonomy: feminist perspectives on autonomy, agency, and the social self*. New York: Oxford University Press.
- Meyers, D. T. (1989). *Self, society, and personal choice*. New York: Columbia University Press.
- Mitchell, G. (2005). Libertarian paternalism is an oxymoron. *Northwestern University Law Review*, 99(3).
- Morozov, E. (2014). *To save everything, click here: the folly of technological solutionism*. New York: Public Affairs.
- Nagel, S. K., Hrinco, V., & Reiner, P. (2016). Algorithm anxiety: do decision-making algorithms pose a threat to autonomy? *IEEE Ethics*, 2016 May 13-14, Vancouver, Canada.
- Owens, J., & Cribb, A. (2017). “My Fitbit Thinks I Can Do Better!” Do health promoting wearable technologies support personal autonomy? *Philosophy & Technology*. <https://doi.org/10.1007/s13347-017-0266-2>.
- Shiffrin, S. (2000). Paternalism, unconscionability doctrine, and accommodation. *Philosophy & Public Affairs*, 29(3), 205-250.
- Smids, J. (2018). *Persuasive technology, allocation of control, and mobility: an ethical analysis*. Eindhoven: Technische Universiteit Eindhoven.
- Spahn, A. (2012). And Lead us (not) into persuasion...? Persuasive technology and the ethics of communication. *Science and Engineering Ethics*, 18, 633–650.
- Specker Sullivan, L., & Niker, F. (2018). Relational autonomy, paternalism, and maternalism. *Ethical Theory and Moral Practice*. <https://doi.org/10.1007/s10677-018-9900-z>.
- Sunstein, C. R., & Thaler, R. (2003). Libertarian paternalism is not an oxymoron. *The University of Chicago Law Review*, 70(4), 1159–1202.
- Thaler, R. H., & Sunstein, C. R. (2009). *Nudge*. Penguin.
- Thaler, R. H., Sunstein, C. R., & Balz, J. P. (2010). Choice architecture. Available at SSRN: <http://ssrn.com/abstract=1583509> or <https://doi.org/10.2139/ssrn.1583509>
- Tripathi, P., People trust Apple more than Google and Facebook *DazeInfo* April 12, 2018. Accessed online at: <https://dazeinfo.com/2018/04/12/apple-google-microsoft-facebook-most-trusted-company/>
- Verbeek, P.-P. (2009). Ambient intelligence and persuasive technology: the blurring boundaries between human and technology. *Neuroethics*, 3, 231–242.
- Wagner, N.-F. (2018). Doing away with the agential bias: agency and patiency in health monitoring applications. *Philosophy & Technology*. <https://doi.org/10.1007/s13347-018-0313-7>.
- Wolf, S. (1997). Happiness and meaning: two aspects of the good life. *Social Philosophy and Policy*, 14(1), 207.
- Zuckerman, E. (2018). Facebook only cares about Facebook. *The Atlantic*. January 27, 2018. Accessed online: <https://www.theatlantic.com/technology/archive/2018/01/facebook-doesnt-care/551684/>