UvA-DARE (Digital Academic Repository)

Optimization of what? For-profit health apps as manipulative digital environments

Sax, M.

DOI

[10.1007/s10676-020-09576-6](https://doi.org/10.1007/s10676-020-09576-6)

Publication date

2021

Document Version

Author accepted manuscript

Published in

Ethics and Information Technology

[Link to publication](#)

Citation for published version (APA):

Sax, M. (2021). Optimization of what? For-profit health apps as manipulative digital environments. *Ethics and Information Technology*, 23(3), 345-361.
<https://doi.org/10.1007/s10676-020-09576-6>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<https://dare.uva.nl>)

Optimization of What? For-Profit Health Apps as Manipulative Digital Environments

Marijn Sax¹

Abstract

Mobile health applications ('health apps') that promise the user to help her with some aspect of her health are very popular: for-profit apps such as MyFitnessPal, Fitbit, or Headspace have tens of millions of users each. For-profit health apps are designed and run as *optimization systems*. One would expect that these health apps aim to optimize the health of the user, but in reality they aim to optimize user *engagement* and, in effect, *conversion*. This is problematic, I argue, because digital health environments that aim to optimize user engagement risk being manipulative. To develop this argument, I first provide a brief analysis of the underlying business models and the resulting designs of the digital environments provided by popular for-profit health apps. In a second step, I present a concept of manipulation that can help analyze digital environments such as health apps. In the last part of the article, I use my concept of manipulation to analyze the manipulative potential of for-profit health apps. Although for-profit health can certainly empower their users, the *conditions for* empowerment also largely overlap with the conditions for manipulation. As a result, we should be cautious when embracing the empowerment discourse surrounding health apps. An additional aim of this article is to contribute to the rapidly growing literature on digital choice architectures and the ethics of influencing behavior through such choice architectures. I take health apps to be a paradigmatic example of digital choice architectures that give rise to ethical questions, so my analysis of the manipulative potential of health apps can also inform the larger literature on digital choice architectures.

¹ Postdoctoral researcher at the Institute for Information Law (IViR), University of Amsterdam.

Keywords: manipulation, health apps, mhealth, autonomy, choice architectures.

The study was supported by the Research Priority Area ‘Personalised Communication’ of the University of Amsterdam

1. Introduction

Mobile health applications (‘health apps’) that promise the user to help her with some aspect of her health are very popular: apps such as MyFitnessPal, Fitbit, or Headspace have tens of millions of users each. Most of the popular health apps are for-profit apps. There is a popular, for-profit health app available for a wide range of health and lifestyle challenges people experience, such as sticking to a diet, living a mindful life, and working out regularly. Everyone with a smartphone can download a health app in a matter of seconds, the apps are designed to be easy to use, and by using the apps of their own choice people can work on their health in the ways they themselves see fit. As a result, health app developers, health care providers, and health insurers alike tout the potential of health apps as *tools of empowerment*. Even public policy strategies as formulated by, for instance, the European Commission (2012, 2014) and the World Health Organization (2011) focus on the great potential of health apps to improve the health of large populations.

Recently, different authors have pointed out how the digital environments we navigate everyday can shape and steer our behavior in significant ways. Zuboff (2015, 2019) has argued that many of these digital environments commodify and attempt to modify our behavior to serve the (financial) interests of the designers of these digital environments, instead of the interests of the users. Frischmann and Selinger (2018: 3) observe “the beginning of a path where powerful companies use smart technologies to gain control over us by framing our choice and nudging us towards programmed lives of convenience and cheap bliss.” Susser (2019: 1) warns that ““adaptive choice architectures” [...] work to subtly guide us toward certain ends.”

In this article, I focus on popular for-profit health apps² as paradigmatic examples of intentionally designed digital environments and argue that they are not exempt from these worries, despite all the positive press they receive. Health apps certainly can have an empowering effect, but the ways in which especially the more popular, for-profit health apps are designed and operated also requires us to critically question this promise of empowerment. The underlying business models of health apps require such apps to be designed and run like *optimization systems*, meaning their “organizing principle is *optimization*” (Overdorf *et al.* 2018: 1). One would expect that health apps are aimed at optimizing *the health* of their users, but in reality the need to monetize their userbase leads for-profit health apps to, first and foremost, optimize user *engagement* and, in effect, *conversion*³ of users. What is particular to the health app context, is that the optimization of user engagement and user conversion can be achieved by framing the functioning of the app in terms of the optimization of the users’ health. Since everyone wants to be healthier, and everyone quite literally needs one’s health to live a fulfilling life, promises of improved health are an interesting commercial strategy (Sax, Helberger & Bol 2018). However, we should at least consider the possibility that the targeting and exploitation of people’s desire for health can lead to problematic influences on app users’ behavior. More specifically, I argue – in line with Susser, Roessler, and Nissenbaum (2019a, 2019b) – that the concept of *manipulation* can be useful in explaining why the influences that digital environments exert on our behavior can be problematic. Popular for-profit health apps, I argue, are a paradigmatic example of a digital environment that risks being *manipulative*.⁴ It is important to gain a better understanding of the risk of manipulation in digital choice environments such as health apps, because the conditions for empowerment and the conditions for manipulation largely overlap in such environments. To support or empower people, digital choice environments can collect and analyze user data to map personal

² I thus do *not* focus on non-profit health apps, or on the more medical health apps that are built – often in cooperation with academic hospitals – to address a very specific medical problem. I focus on the big commercial players that develop ‘healthy lifestyle and wellness’ services for the general consumer population.

³ Conversion refers to the principle of turning users into profitable users.

⁴ Because of the limited space at my disposal, I will not explicitly engage with related challenges posed by health apps, such as challenges of informational privacy (Floridi 2005, Patterson 2013, Lanzing 2016) and surveillance (Lupton 2012).

characteristics and circumstances which can be used to personalize services in a manner that will respond to the needs and desires of individual users. However, that same access to information about people, as well as the access to their decisional sphere, can also be *misused* for manipulative strategies which target personal characteristics or circumstances to (try to) steer people's behavior in the direction the manipulator desires. In practice, it is and will remain difficult to draw a sharp line between (conditions for) empowerment and (conditions for) manipulation.

In response to this challenge, this article has two distinct but related aims: it seeks to 1) further our understanding of the concept of manipulation and how it can be applied to digital environments that attempt to influence their users; and 2) to nuance the empowerment discourse around health apps by showing how the digital environments they provide to users can also have manipulative tendencies.

This article is structured as follows. The first part analyzes the way for-profit health apps operate and the digital environments they provide to users. Here I discuss the underlying business models of health apps, how those business models help explain why health apps are designed and run as optimization systems aimed at optimizing engagement and conversion, and how data analytics plays an important role in enabling such optimization. Besides data analytics, health apps also deploy a particular health discourse which is instrumental to their business practices. Second, I turn to the concept of manipulation. This often-used but under-analyzed concept needs clarification before it can be used to analyze digital health app environments such as health apps. I largely follow Susser, Roessler, and Nissenbaum (2019a, 2019b) and argue that manipulation should be understood as the intentional infiltration of the decision-making of a manipulee with the aim of making the manipulee a pawn in the self-serving scheme of the manipulator. In doing so, a manipulator disregards, or cleverly encapsulates (Hardin 2002),⁵ the true interests of the manipulee. Manipulators target known, presumed or inferred weaknesses in their target in an attempt to 'push their buttons' or 'pull their strings.' Lastly, manipulators attempt to keep their manipulation under the radar, although the success of manipulation does not necessarily depend on the manipulative influences remaining

⁵ Thanks to Thomas Nys for pointing this out to me. See also Nys 2016.

completely hidden. Here I also briefly explain why manipulation is problematic, by discussing how it undermines personal autonomy. Third, I turn to health apps again to analyze how the digital environments they provide to users risk being manipulative and how their manipulative potential can be difficult to disentangle from their potential to support or empower their users.

2. Apps as Optimization Systems

We need to understand how for-profit health apps operate before we can analyze the potentially manipulative character of the digital environments they provide to users. I suggest we start with the underlying business models of for-profit health apps. Above all, for-profit health apps need to monetize their userbase to be successful. To understand how they monetize their userbase, we need to analyze their business models. As Apple explains on its developer platform⁶ which is aimed at helping app developers build successful apps: “Consider choosing your business model before you start developing your app so you can build it seamlessly into the user experience.”⁷

It is, of course, entirely unsurprising that for-profit health apps seek to monetize their userbase to realize their profit-seeking ends. I want to emphasize at the outset that my argument is *not* concerned with the profit-driven motives of health apps per se. Rather, my argument is that the need to monetize their userbase incentivizes for-profit health apps to organize their app around the principles of engagement and conversion optimization. In the health app context, the pursuit of those optimization principles can – but does not necessarily have to – lead to manipulative digital environments.

So let us start with a brief analysis of business models for mobile apps. On their developer platforms, both Apple and Google provide a helpful overview of different business models.⁸ (Notice that most of the following business models are not mutually exclusive.)

⁶ See Fahy *et al.* (2018) for a more elaborate analysis of Apple’s and Google’s developer platforms and the different kinds of monetization strategies apps can use.

⁷ Apple, ‘Choosing a Business Model’: <https://developer.apple.com/app-store/business-models/>

⁸ Apple, ‘Choosing a Business Model’: <https://developer.apple.com/app-store/business-models/>.

Google, ‘Earn more revenue with the right monetization options’: <https://developer.android.com/distribute/best-practices/earn/monetization-options>.

- Users pay a one-time fee to install and use an app.
- Selling space for advertising. Google helpfully suggests to consider “native ads that allow you to match ads to your app’s look and feel.”⁹ Native advertising is “also called sponsored content [...] a term to describe any paid advertising that takes the specific form and appearance of editorial content from the publisher itself” (Wojdyski & Evans 2016: 157).¹⁰
- Offer in-app purchases, either in the form of physical products, or in the form of additional content or features.
- Users pay recurring subscription fees in order to gain or retain access to either the entire app, or to some additional content or features.
- Monetization of user data.¹¹

Although different health apps are built on the basis of different business models, all business models give rise to the same incentive: making sure one has as many active users as possible. Whether one’s income depends on letting users pay for access, showing (native) advertising to users, offering in-app purchase options to users, selling subscriptions, or monetizing user data, in all cases more active users means more potential to generate income. To achieve this, app designers try to build digital environments that optimize for user *engagement*: “The path to monetization is through engagement, and when users are given time to enjoy an app, they may be more inclined to invest in paid features.”¹² Simply put, app providers want users to use an app for an extended period of time, spend as much time as possible in the app, and engage with as much revenue-generating features (e.g., premium features) and material (e.g., (native) advertising) as possible. Because app providers are in charge of the digital environment *within which* the

⁹ Google, ‘Earn more revenue with the right monetization options’: <https://developer.android.com/distribute/best-practices/earn/monetization-options>

¹⁰ In the case of health apps this means advertising that is deliberately designed to ‘look and feel’ like health content, thereby obfuscating to the user the true commercial nature of the content.

¹¹ This option is not explicitly mentioned by Apple or Google. It is, however, a serious option. MyFitnessPal, an immensely popular calorie counting app with millions of users, is a good example. On its jobs page, MyFitnessPal announces that “MyFitnessPal has the largest database of human eating habits in the world. The opportunities for a data scientist here are almost endless.” (<https://www.myfitnesspal.com/jobs>) Forbes reports that health care providers and researchers can access the database when they enter into a “formal partnership” with MyFitnessPal (Olson 2014).

¹² Apple, ‘Using the Freemium Model’: <https://developer.apple.com/app-store/freemium-business-model/>

relationship between user and app develops, they will attempt to design the digital environment in such a manner that the engagement of users with the app is optimized. To know *how* the digital environment can be optimized for engagement, data analytics plays an essential role. With the right tools, engagement can be *engineered* (see, for instance, Eyal 2014 and Alter 2017).

Various tools for data analytics are available and easily accessible to app developers. For example, Google's platform offers built-in analytics tools to make things easier for app developers. Google offers 'Google Analytics for Firebase,' which can be used to gain a detailed understanding of how users behave *within* an app by tracking "500 distinct events."¹³ In its best practices for Firebase, Google explains that the analytics tool can be used to, for instance, "Dynamically tailor your app's features to specific audiences. Use Remote Config to change the look and feel of your app for a specific audience."¹⁴ Another possibility is to "Improve your acquisition workflow. Use the integration with Google Ads to understand the influence of your advertising and marketing activities. Ensure your campaigns are acquiring engaged and valuable users by tracking the app open events and automatically linking user behavior within your app to a traffic source."¹⁵ Through its built-in tools, Google thus offers tremendous tracking, analytical, and targeting potential to app developers. Apps can be run like large scale, never-ending experiments to test how the app must be designed to make sure users keep coming back, spend as much time as possible in the app, and engage with as much revenue-generating features and material as possible. Through feedback loops, the digital environment can even be personalized in real time, to truly optimize for engagement (Yeung 2017, Lanzing 2018). By optimizing for user engagement, apps automatically optimize for conversion – i.e., the conversion of users into profitable users – because engaged users interact with revenue-generating features and material in the apps.

¹³ Google, 'Improve Conversion Using Google Analytics for Firebase': <https://developer.android.com/distribute/best-practices/earn/improve-conversions>.
¹⁴ Google, 'Improve Conversion Using Google Analytics for Firebase': <https://developer.android.com/distribute/best-practices/earn/improve-conversions>.
¹⁵ Google, 'Improve Conversion Using Google Analytics for Firebase': <https://developer.android.com/distribute/best-practices/earn/improve-conversions>.

2.1. An Example of Optimization in Health Apps

To see how this general description of the app economy holds for health apps, let us briefly discuss one example. The health app Headspace is a good example of a health app that communicates to its users the aim of optimizing their (mental) health, while the design and operation of the app betray a primary focus on optimizing engagement and conversion. Headspace is a very popular meditation and mindfulness app with tens of millions of users that offers a range of meditation exercises aimed “to improve the health and happiness of the world.”¹⁶ These exercises are organized in ‘packs,’ which contain multiple ‘sessions,’ and ‘singles,’ which are stand-alone meditation sessions. Headspace is a typical freemium app, with a business model based on luring people to the app with a few free-to-use packs containing basic meditation sessions. Users can then buy additional themed singles (such as ‘commuting,’ and ‘walking in the city,’) and packs (such as the ‘Health packs’ named ‘Coping with Cravings,’ and ‘Coping with Cancer’), or they can choose for monthly, yearly or even eternal subscriptions that give access to all available content. Headspace does not offer any additional products or services, so everything is built around the different meditation sessions. The layout of the app is very clean, friendly, cheerful, and basic. Very few information is shown on screen and they make use of soft, soothing colors. Put shortly, Headspace looks and feels like a very simple and low tech app that just offers some meditation exercises.

Precisely because Headspace is such a simple and unobtrusive app that claims to merely aim to make its users happier and healthier, and does not feel like a technologically advanced service at all, it is an interesting test case. Health apps like Headspace are rather secretive when it comes to their business practices which they treat as trade secrets securing their competitive edge. As a result, finding definitive ‘proof’ of, or ‘evidence’ for, the specific practices that Headspaces engages in is almost impossible.¹⁷ As a workaround, I looked into the advertised job openings which mention specific business goals and request specific skills and knowledge. The job openings at Headspace suggest that the app is in fact

¹⁶ Headspace, ‘About Headspace’: <https://www.headspace.com/about-us>.

¹⁷ This lack of access to the (business) operations of Big Tech companies could of course be criticized for a variety of reasons. For example, it makes it harder for investigative journalists and academics to scrutinize the practices of these companies. The same holds for policymakers and regulators who often have a hard time gaining access to Big Tech companies.

aiming to, above all, optimize engagement (and in effect conversion) of users by employing advanced data analytics. In October 2018, Headspace was looking for a new Senior Data Analyst that can “[d]rive projects to identify key levers for new user growth, retention, and revenue” and “[l]everage data to understand the product, identify opportunities, and execute initiatives to drive growth and engagement” and “[d]rive experiment design, interpretation, and actionable insights.”¹⁸ Another example is the job opening for a new VP of Growth: “you will be responsible for growing and engaging the Headspace audience for all our Acquisition, CRM, Conversion, and Monetization efforts. You will have a bias towards experimentation and innovation to continue to identify new acquisition channels, while optimizing conversion and retention.”¹⁹ The ideal candidate is “Curious – you are hyper-focused on data quality and are data driven, you are always looking [...] to find new way to attract and retain customers.”²⁰ Such job descriptions suggest (but do not prove), that even the simplest, most unobtrusive, friendly looking mindfulness and meditation health app might be run as an engagement and conversion optimization system with the help of advanced tools for analytics. Because the work that employees of Headspace do might deviate from the activities and objectives described in the job descriptions, we cannot be certain that the information provided in Headspace’s job openings is an accurate reflection of the internal operations of the company. When Forbes visited Headspace they came to the same conclusion. “But perhaps Headspace’s most important strength lies in data, which it has been using to understand what makes new users become regulars, when people zone out of tracks and how the app can become personalized to predict users’ needs” (Chaykowski 2017).

To be able to analyze to what extent the digital environments provided by health apps are problematic – if at all – we also need to briefly look into the health discourse found in health apps. Though the propagation of a rather particular health discourse, health apps seek to influence their users’ understanding of their (relation to their own) health. The

¹⁸ Headspace, ‘Senior Data Analyst,’ job description that has since been removed, screenshot available here: <https://imgur.com/a/qtSe4li>.

¹⁹ Headspace, ‘VP of Growth,’ job description that has since been removed, screenshot available here: <https://imgur.com/a/qtSe4li>.

²⁰ Headspace, ‘VP of Growth,’ job description that has since been removed, screenshot available here: <https://imgur.com/a/qtSe4li>.

remainder of this section discusses health discourse. The next section develops an account of manipulation that can be used to analyze the digital health app environments.

2.2. *Strategic Health Discourse: Health-as-Wellness*

Health is a contested concept. Philosophically, both the meaning and value of health are subject to debates (see, for instance, the influential debate between Boorse (1975, 1977) and Nordenfelt (1986, 1987).) In everyday life, the term ‘health’ is also used to refer to a wide range of practices, products, and services, rendering the precise meaning of the term ‘health’ unclear. Health apps can exploit this conceptual fluidity to propagate a particular understanding of health.²¹

Health apps want to use the term ‘health’ because health is a universally desired and needed good. Health, like sex, sells, and using the term can *entice* people to use a (self-proclaimed) health app, or to engage with content, products, or service labeled as ‘relevant to health’ *within* an app. Because using a health frame is so attractive to health app providers, they do not want to be stuck with an overly medicalized and narrow notion of health since that would inhibit their ability to frame services and products in terms of health. The way health apps secure the benefits of using the term health with all its positive connotations, while still being able to frame nearly everything as health-related, is to collapse the concept of health into the broader and vaguer concept of *wellness* (Cederström & Spicer 2015). Wellness refers to the general idea of ‘doing well’ and ‘feeling good,’ which, in turn, can refer to nearly anything and everything in a person’s life. Brodesser-Akner (2018) accurately captures the contemporary catch-all nature of the health-as-wellness discourse:

Before we knew it, the wellness point of view had invaded everything in our lives:
Summer-solstice sales are wellness. Yoga in the park is wellness. Yoga at work is
wellness. Yoga in Times Square is peak wellness. When people give you namaste

²¹ There is a rich literature on Foucauldian biopower and health and the role (digital) technologies can play in the exercise of biopower (e.g., Foucault 1975, Armstrong 1995, Petersen & Bunton 1997, Casper & Morrison 2010, Lupton 2012, Mayes 2015, Ajana 2017, Fotopoulou & O’Riordan 2017, Sanders 2017). Although this literature provides interesting and promising perspectives for my research, I do not have enough space at my disposal in this article to incorporate this complex literature into my argument.

hands and bow as a way of saying thank you. The organic produce section of Whole Foods. Whole Foods. Hemp. Oprah. CBD. "Body work." Reiki. So is: SoulCycle, açai, antioxidants, the phrase "mind-body," meditation, the mindfulness jar my son brought home from school, kombucha, chai, juice bars, oat milk, almond milk, all the milks from substances that can't technically be milked, clean anything. "Living your best life." "Living your truth." Crystals.

What emerges is 'health-as-wellness:' a catch-all conception of health. Online advertising agency *NativeAdBuzz* posted a New Year's resolutions message in December 2015 that exemplifies the strategic use of the health-as-wellness perspective:

The extreme desire for health and wellness that's been building over the past few years and is about to EXPLODE across the web (while sprinkling billions of dollars in profits out to various publishers and affiliates across 100+ different countries)...

Are you going to be one of the publishers or advertisers who reaches out and grabs a big piece of the health and fitness dough that's openly available for the taking?²²

Cederström and Spicer (2015: 3, 5) also observe that wellness has become a norm that people should adhere to: "Today wellness is not something we choose. It is a moral obligation. [...] As consumers, we are required to curate a lifestyle aimed at maximizing our wellbeing [...] wellness has wormed itself into every aspect of our lives." The other side of the *demand* to adopt a lifestyle aimed at wellness is the *responsibility* of people to actually adhere to this demand. Lupton (2013: 397) explains that these demands and the following responsabilization give rise to 'healthism:' "Healthism positions the achievement and maintenance of good health above many other aspects of life and features of one's identity, so that an individual's everyday activities and thoughts are continually directed towards this goal."

²² NativeAdBuzz, "This Health and Wellness Boom Has Been Building for Years... And It's Finally About to ERUPT (Urgent: Your Free VIP Christmas Gift Has Arrived)": <http://www.nativeadbuzz.com/blog/this-health-and-wellness-boom-has-been-building-for-years-and-its-finally-about-to-erupt-urgent-your-free-vip-christmas-gift-has-arrived/>

When we look at popular for-profit health apps, we also encounter an insistence on the importance of being preoccupied with one's health in the health-as-wellness-sense of the word. Take, for instance, Fitbit. On its 'Why Fitbit' page, one can read that "Every moment matters and every bit makes a big impact. Because fitness is the sum of your life. That's the idea Fitbit was built on—that fitness is not just about gym time. It's all the time."²³ Another example is MyFitnessPal. The app is advertised as a simple calorie counter, but the actual user experience is structured around the MyFitnessPal blog which takes up around 80% of the home screen. With multiple post a day, the blog constantly reminds the user what a perfectly curated healthy lifestyle looks like. It does so by framing a wide range of activities and products as relevant to the user's health: the music playlists one listens to²⁴ and the headphones one uses,²⁵ the alarm clock one uses,²⁶ the skincare products one's yoga instructor uses,²⁷ and how one spends \$50 at Whole Foods.²⁸

In popular for-profit health apps, health is thus portrayed as something that is about everything you do, and as something you should be preoccupied with. The health apps conveniently provide you with products, services, and subscriptions that help you in your pursuit of the curated healthy lifestyle that is presented to you in the very same app.

In the first part of this article, we have seen that popular for-profit health apps can be understood as optimization systems that, first and foremost, optimize user engagement and conversion. They do so by building digital environments that allow them to conduct experiments continuously in order to find the most effective ways of targeting users and personalizing the app experience. Health apps also rely on a health discourse that allows them to frame nearly every practice, service, or product as relevant to the user's health, while at the same time communicating to the user the importance of being preoccupied

²³ Fitbit, 'Why Fitbit': <https://www.fitbit.com/whyfitbit>.

²⁴ MyFitnessPal, 'These Playlists Were Built to Make You Better': <https://blog.myfitnesspal.com/these-playlists-were-built-to-make-you-better/>

²⁵ MyFitnessPal, 'These On-Ear Headphones Can Actually Withstand Your Workouts': <https://blog.myfitnesspal.com/these-on-ear-headphones-can-actually-withstand-your-workouts/>

²⁶ MyFitnessPal, 'Why and How You Should Nix an Alarm Clock': <https://blog.myfitnesspal.com/why-and-how-you-should-nix-an-alarm-clock/>

²⁷ MyFitnessPal, 'A Day in the Life of a Yoga Teacher': <https://blog.myfitnesspal.com/day-life-yoga-teacher/>

²⁸ MyFitnessPal, 'How a Nutritionist Spends \$50 at Whole Foods': <https://blog.myfitnesspal.com/how-a-nutritionist-spends-50-at-whole-foods/>

with one's health. In the next section, I develop a conception of manipulation (largely in line with Susser, Roessler and Nissenbaum 2019a, 2019b) to help analyze whether – and if so, why – the digital environments that for-profit health apps provide to their users are problematic. The guiding intuition is that popular for-profit health apps target and to a certain extent *misuse* people's natural desire for health to get them to spend more time and possibly more money on health apps than they, all things considered, would want to.

3. Manipulation

In everyday life, the term 'manipulation' is used to describe a wide range of phenomena. Noggle (1996: 43) even speaks of the "commonness of manipulation in everyday life." For example, people say that they manipulate *objects* such as levers and juggling balls when they move them around at their will. People also say that they manipulate *persons* such as their neighbors, friends, and lovers. Lastly, it is said that *more abstract entities* such as institutions are manipulated (e.g., elections), or that we ourselves are manipulated by them (e.g., by the dating app Tinder which sells premium 'pay-to-win' features to users by exploiting their fear of being put at a disadvantage vis-à-vis their dating competitors who are already using premium features to boost their chances). Precisely because the term is used in so many different situations, formulating a sufficiently coherent conception of manipulation is challenging.

Before we move on and try to actually define manipulation, we should therefore ask ourselves what purpose we would like a definition to serve. One option is to look for a definition that tracks as precisely as possible all the colloquial uses of the term. Another option is to define manipulation as a moralized concept which captures a particular kind of problematic influence that warrants special attention. In what follows, I argue that manipulation is a form of influence that should be treated as a moral wrong²⁹ because it instrumentalizes persons and it does not respect their autonomy.

First, I will develop my conception of manipulation and briefly discuss how manipulation undermines personal autonomy. Second, I will address two questions my concept of manipulation may give rise to.

²⁹ Contrary to authors like Buss (2005) and Wood (2014) who treat manipulation as a non-moralized term.

3.1. Conceptualizing Manipulation

A core feature of cases of interpersonal manipulation is that manipulators *use* manipulees as pawns in their schemes. Seen from this perspective, when one person manipulates another person, she sees and uses her as if she were an object that can simply be used as the manipulator sees fit – manipulation *instrumentalizes* the manipulee. Noggle (1996: 44), for instance, write that “The term “manipulation” suggests that the victim is treated as though she were some sort of object or machine.”

The analogy to objects, however, also immediately shows an important difference between manipulating objects and persons. Since objects are inanimate, we can simply move (manipulate) them precisely as we see fit: they do not complain and have no interests. Persons are, of course, different. They have to be ‘steered’ via some form of influence for them to become a useful pawn in the manipulator’s scheme. It is therefore not sufficient to say that in cases of manipulation one person uses another person as a pawn in her scheme. Consider one of those classic movie scenes where a bandit wants to stop a train in the middle of the desert by leaving a tied up person on the tracks in clear sight. When the train operator sees the person on the tracks, she is forced to stop the train in order to save the tied up person’s life. The bandit is surely using the tied up person as a pawn in her scheme, but we would not say that the bandit is *manipulating* the tied up person. Manipulation, then, is associated with particular *forms* of exerting influences on persons in order to use them for our own ends.

So what forms of influence would constitute manipulation? There exist many different ways to influence persons, ranging from outright coercion to purely rational persuasion. Manipulation would be something in between these two extremes. So let us first look at these two extreme ends of the ‘influence scale’ in order to gain an understanding of what manipulation is *not*. As a result, we can better understand what sets manipulation apart from other paradigmatic forms of influence.

Rational persuasion works by offering another person *reasons* (often in the form of arguments) for doing or wanting something. What makes persuasion *rational* is the fact that the offering of reasons happens in a transparent fashion and appeals to the rationality of the other person; ideally no additional influences – other than arguments that appeal to

reason – are used to persuade. It is then up to the other person to freely deliberate – either internally or with the persuader or other persons – about these reasons.

We speak of coercion when the coercer puts another person in a position in which the other person cannot *reasonably* do anything other than comply with the coercer's demands. Although the coerced person is – formally speaking – able to choose otherwise, the coercer has made other options so unattractive that complying with the coercer's demands is the only real option. Hayek (2006 [1960]: 89), for instance, writes that “Though the coerced still chooses, the alternatives are determined for him by the coercer so that he will choose what the coercer wants.”

Both rational persuasion and coercion can be effective means of influencing another person, but they lack something which seems to be central to manipulation. Both rational persuasion and coercion are very *straightforward* ways to influence someone because one either explicitly engages with someone in the form of arguments and reasons (persuasion) or makes it abundantly clear which of the alternatives someone should choose (coercion). Manipulation, to the contrary, is a more roundabout and “subtle and sneaky” form of influence where “the manipulator *infiltrates* their decision-making process, disposing it to the manipulator's ends, which may or may not match our own” (Susser, Roessler & Nissenbaum 2019a: 17, emphasis mine). Starting from this characterization of manipulation, I will discuss what I take to be the necessary features of manipulation.

First, a manipulator infiltrates the decision-making of another person precisely because the manipulator wants to further her own ends by making use of the other person. A manipulator is concerned with her own ends and tries to find ways to get others to serve her ends. In doing so, the true interests and desires of the manipulees do not figure prominently in her scheming. Cases of manipulation are characterized, then, by the manipulator's disregard for, or indifference to, the manipulee's true interests. Notice, however, that the manipulator's disregard for a manipulee's true interests *does not necessarily imply* that the manipulee's interests will not figure in the manipulator's practical reasoning. A manipulator can attempt to ‘encapsulate’ (Hardin 2002, Nys 2016) a manipulee's interests for strategic reasons. But a manipulator will only do this precisely *because* doing so serves her own ends, not because she is genuinely concerned with the manipulee. So even if a manipulator deliberately encapsulates a manipulee's interests, this

could still be characterized as a disregard for the true interests of the manipulee. It follows that manipulators do not need to have as their explicit aim to *harm* their manipulees, although they certainly can harm manipulees in the process of their manipulation. This feature of my concept of manipulation has implications for its relation to everyday uses of the term. In everyday life, people often say that they manipulated their loved ones, for instance in a 'playful' manner ("Gerry and I hadn't had sex in quite some time, so I manipulated him to do it") or 'for their own good' ("My kid Beth wanted to go out with her friends again, but I manipulated her into staying home to work on her homework".) In such cases, people engage in their self-proclaimed acts of manipulation precisely because they do care for the (interests and desires of the) other person. My moralized conception of manipulation does not account for such cases. I find this implication a strength rather than a weakness, since I do not consider the so-called 'manipulation' of loved ones a moral wrong.

Second, because manipulators resort to manipulation to dispose us to their ends, manipulation is by definition intentional on the part of the manipulator. Manipulators manipulate deliberately to get things done. It would thus be conceptually incoherent to claim that someone was manipulated accidentally. When reflecting on an event, a person can certainly *feel* manipulated irrespective of whether the other person(s) involved *intended* to manipulate the person. But we would only speak of manipulation when the manipulator *intended* to make the manipulee into a pawn in her scheme.

Third, a manipulator attempts to *infiltrate* someone's decision-making process – rather than engaging with it in a straightforward manner – because the manipulator is reasonably sure that her target is not willing to (fully) cooperate if asked. So a manipulator will need to find out which 'buttons to push' or which 'strings to pull'; manipulators ideally know what makes us tick. Manipulators, then, will attempt to identify *any* characteristic³⁰ of a person's psychology that can be exploited to steer that person's behavior in the (self-serving) directions that the manipulator desires. In the context of manipulation,

³⁰ Susser, Roessler, and Nissenbaum (2019a) speak of 'vulnerabilities.' I rather use the looser term 'exploitable characteristics of a person' because the term 'vulnerabilities' is sometimes associated (especially in legal discourse) with a fixed set of narrowly defined weakness, such as those that are the result of one's age ('the old and the young') or of one's psychical or mental infirmities ('people with medically diagnosed handicaps.')

information often equals power. The more one knows about the potential target, the better one will generally be able to identify ‘exploitable characteristics.’³¹ In the absence of perfect information, manipulators will target presumed or inferred exploitable characteristics. Despite the fact that every person is unique in some way, people are also conveniently similar in many other respects. Every person has basic needs (such as the need for love and health) which can be targeted by manipulators. The behavioral-economics literature that was popularized by Thaler and Sunstein (2008) and Kahneman (2011) also provides a rich source of features of the human psychology that can be targeted by manipulators.

Fourth and last, manipulators will never explicitly announce or draw attention to the fact that they are attempting to manipulate someone. Manipulation typically works best when the manipulees are either unaware of the fact that someone is (trying to) push their buttons, or are unaware of the techniques that are used to push their buttons. It is, however, possible that a manipulator successfully disposes a manipulee to her ends *even though* the manipulee becomes aware of (the working of) the manipulative influence.³² Manipulators always attempt to target exploitable characteristics of a person and if such targeting is done well enough, the manipulee can still feel compelled to act as the manipulator intended, without the manipulative influence raising to the level of outright coercion.

In sum, I argue that manipulation is an infiltration of decision-making that is (1) intentional, (2) seeks to further the interests or ends of the manipulator by making use of the manipulee, while disregarding or encapsulating the true interests of the manipulee, (3) happens through the targeting of presumed, known, or inferred exploitable characteristics of the manipulee, and (4) is never announced or emphasized by the manipulator (even though manipulation may still occur when the manipulee discovers the attempt.)

³¹ Rudinow (1978: 346, emphasis added) explains that “the manipulator’s behavior is normally either deceptive or predicated on some *privileged insight into the personality of his intended manipulee*.”

³² Mills (2014: 138) provides a similar argument, referring to Gorin (2014) and Barnhill (2014): “Both Gorin and Barnhill point out that manipulation does not need to involve deception or covertness; these are not defining features of manipulation necessarily present in all cases of what we could agree to be manipulation. But most manipulators seek to hide the degree to which they are angling to achieve their desired result and would find the success of their project seriously compromised if their manipulative intentions were revealed.”

3.2. *Manipulation and Autonomy*

At the beginning of this section, I explained that I use a moralized conception of manipulation, meaning that on my account manipulation is a moral wrong. It might be suggested that manipulation is wrong because manipulees may end up acting against their own interest. However, as Susser, Roessler, and Nissenbaum (2019a: 35, emphasis added) rightly point out, “beyond the *material harms* that result from manipulation, such as exploitation, impoverishment, unfairness, and the deprivation of benefits, the deeper harm is infringement individual autonomy.” I do not have the space at my disposal to offer an elaborate account of autonomy, so below I will only indicate briefly how manipulation threatens autonomy.³³

Put simply, an autonomous person is a self-determining person that is able to live her life in self-chosen ways. Standard procedural accounts of autonomy specify this ideal by emphasizing the importance of “*competence* in reflection and decision making and (on some views) *authenticity* of values, desires, and so on that constitute the person and motivate choice.” (Christman 2004: 148, emphasis added) On such an account, an autonomous person can deliberate about her own situation – which involves, on the one hand, her desires, goals, and values, and, on the other hand, the available information and the material conditions she finds herself in – and decide for herself how she would like to act. Moreover, it is required that she is able to critically reflect on her own desires, goals, and values, to check whether these are *authentic* – i.e., whether they are genuinely her own.³⁴ If proper deliberation which involves authentic desires, goals, and values has taken place, then the ‘output’ of this process can be considered an autonomous decision.

On such a standard procedural account of autonomy, manipulation fails to respect autonomy and threatens the practicing of it. Manipulation is characterized by the instrumentalization of manipulees; they are treated as if they were mere instruments to be used by the manipulator instead of people capable of deciding and living autonomously. Moreover, when manipulators successfully make manipulees serve their interests by

³³ Susser, Roessler, and Nissenbaum (2019a, 2019b) provide a more elaborate discussion on the connection between autonomy and manipulation.

³⁴ Christman (1991: 10) has formulated a popular, somewhat weaker alternative: “What matters is what the agent thinks about the process of coming to have the desire, and whether she resists that process when (or if) given the chance.”

targeting exploitable characteristics, they are bypassing the manipulees' capacity to deliberate in an autonomous fashion.

In the contemporary literature, autonomy's relational nature has been emphasized by a range of authors (e.g., Meyers 1989, Nedelsky 1989, Code 1991, Mackenzie and Stoljar 2000, Anderson and Honneth 2005, Stoljar 2011). Standard procedural accounts of autonomy are argued to be too individualistic because they both neglect the importance of personal relations in fostering a person's autonomy and seem to assume that a person should, ideally, strive to be as independent and self-determining as possible. Instead of simply rejecting the ideal of self-determination, however, accounts of relational autonomy argue that more attention should be paid to the ways in which our social relations are "causally necessary for and even constitutive of autonomy" (Nys 2016: 11). Anderson and Honneth (2005), for instance, argue that we need to understand ourselves to be person capable of practicing autonomy, before we can actually do so. That self-understanding can only be developed and sustained under "socially supportive conditions" where other people *recognize* us as autonomous persons (Anderson and Honneth 2015: 130).

Given such a more relational perspective on autonomy, manipulation can undermine autonomy as well. A manipulator is concerned with *using* another person to serve her own interest. In doing so, a manipulator is certainly not aiming to provide the manipulee with "socially supportive" conditions to help her understand herself as an autonomous person that can (and should) act on the basis of her own reasons. If anything, it is in a manipulator's self-serving interest *not* to be supportive of the target's autonomy.

3.3. Two Possible Objections: The Means of Manipulation and Hiddenness

In this section I discuss two possible objections to my conception of manipulation. First, it could be argued that my account fails to identify particular *means of manipulation*, which makes the account unclear. Second, it could be argued that manipulation necessarily operates 'in the dark' or 'behind the back' of its targets and that my accounts fails to address this.

3.3.1. Means of Manipulation

Let me start with the question about the *means of manipulation*. Most of the philosophical treatments of manipulation start by exploring (often to later refute) the possibility that manipulation is, necessarily, a form of deception (see, e.g., Rudinow 1978, Noggle 1996, Baron 2003, Greenspan 2003, Cohen 2018, Susser, Roessler & Nissenbaum 2019a). The appeal of a deception-based definition of manipulation is clear. In many cases that we would intuitively call manipulation, the manipulator resorts to manipulation because she seeks to secure cooperation of the manipulee without wanting to ask for it directly. To still secure the manipulee's usefulness to the manipulator, the manipulee must somehow be "led astray" (Noggle 1996: 44) so that she unknowingly and/or unwillingly becomes a pawn in the manipulator's scheme. Deception seems like a very suitable technique.

I would argue, however, that a real manipulator does not devise her schemes by starting with a fixed set of neatly circumscribed 'manipulation techniques,' but rather starts with the ends and looks for means that can help her achieve those ends. Consider, for instance, the many techniques digital choice architects have at their disposal (see section 2) to tweak the digital environments they offer to people. It should be noted, however, that some techniques to influence or steer behavior – such as outright coercion – are inherently non-manipulative. So what typifies a manipulator is a mindset of seeing others as instruments that can be used to effectuate the desired outcomes. But in the process of looking for means to steer behavior, it can happen that a person with a mindset that could be seen as manipulative, settles on the use of techniques that do not qualify as manipulation as I have defined it.³⁵

3.3.2. *Hiddenness of Manipulation*

Another question that could be raised about my account is why it does not require manipulation to be hidden to the manipulee. Susser, Roessler and Nissenbaum (2019a)

³⁵ I admit that the resulting picture can feel a bit messy or unclear. Someone can have a manipulative mindset, but, in the end, be drawn to techniques that are *not* manipulative in nature – An in principle manipulative mindset does not *necessarily* lead to manipulation. Although I agree that the resulting picture is messy, I see no way to avoid this. The empirical reality of data-driven dynamically adjustable choice architectures simply *is* very messy. The industry is constantly running (multiple, parallel) experiments to test a plethora of tweaks to their digital choice architectures to test whether those tweaks can successfully shape (patterns of) behavior. In this constant hunt for new behavior influence techniques, some will turn out to be manipulative, while some will turn out to be something else (e.g. coercive).

have recently argued that manipulation is by definition hidden to the manipulee because only a hidden influence effectively steers a person's decision-making while also *alienating* a person from her own decision-making process by robbing her of authorship over her decisions. Susser, Roessler, and Nissenbaum also see manipulation as a moralized concept. Their definition includes the hiddenness of the manipulative influence as a necessary feature of manipulation, because, as I interpret them, they take the 'alienation from one's own decision-making process' to be the most significant moral harm associated with manipulation. Because I think that a slightly different moral harm should be captured by my moralized conception of manipulation, I end up with a slightly different stance on hiddenness as a requirement for manipulation. Let me first briefly discuss their argument, before presenting my own stance in reaction to theirs. They write that

[...] as soon as we become conscious of outside influence, of someone else's plans and how we are implicated in them, we incorporate that influence into our own decision-making. Once you know someone else is trying to get you to do something, that fact becomes a regular part of how you make up your mind. It becomes one of the reasons that helps you explain your actions to yourself. Since we are never totally free of outside influence, what gives us (part) authorship over our own actions is that we regard our own reasons for acting as authoritative. Manipulation thwarts that (Susser, Roessler, Nissenbaum 2019a: 20).

So let us imagine a person who has worked at the data science department of Headspace and knows all there is to know about their attempts to understand and influences their users' behavior. To this person, the existence and (let us assume) working of the Headspace's techniques to target exploitable characteristics is not hidden. Still, it seems implausible to suggest that this person is completely 'unmanipulable.' Even if you are (vaguely) aware of the fact that your buttons are being pushed by Headspace, Headspace can still, at least in some instances, effectively steer your behavior as long as they have successfully figured out which of your buttons are especially sensitive under which conditions.

Susser, Roessler, and Nissenbaum (2019a) would *not* argue that our imaginative Headspace user's behavior cannot be successfully steered by Headspace. Rather, they would argue that Headspace's influence simply stops being manipulation and turns into something else as soon as the influence stops being completely hidden to this person. Their argument thus seems to be based on the normative judgment that the normatively *most* significant feature of the described situation is the fact that the person is (vaguely) aware of the existence and/or inner workings of the influence.

I disagree that the normatively *most* significant feature of the situation of our Headspace user is that she at least knows *how* the cleverly targeted pushing of her buttons disposed her to Headspace's ends. She might still have (part) authorship over her own actions because she is able to explain that "Headspace correctly identified that I am insecure about aspect X of my (mental) health, which means that Headspace will infer that they can best target X at time Y, to try to make me subscribe to vaguely related wellness service Z." In the end, however, we are still left with a situation where Headspace is willing to intentionally develop a system aimed at systematically identifying and targeting exploitable characteristics – often related to (mental) health – in order to sell as many premium packages as possible. The core intuition that I want my concept of manipulation to capture, is the intuition that there is something wrong with designing digital environments which systematically seek to sniff out and target exploitable characteristics of users, and to use those insights to try to make the users serve the interests of the provider of the digital environment, while disregarding or at best encapsulating the interests of the users. It can certainly *help* to understand (vaguely) how and why our buttons are pushed successfully by others (so we can retain some 'authorship'), but in the end we simply do not want to be subjected to systems that are designed to identify continuously how every individuals' buttons can be pushed as effectively as possible in order to make them useful pawn in someone else's scheme. That is the core intuition I'd like to account for with my conception of manipulation and the 'hiddenness condition' is not required for that purpose. I do however acknowledge that manipulation typically works "best in the dark" (Bovens 2009) and that manipulators will *attempt* to keep their manipulation hidden.

3.4. *One Further Challenge: Manipulation as a Success Concept*

I need to address one further challenge before I can explain to what extent popular for-profit health apps are manipulative. As Faden and Beauchamp (1986: 354), and Susser, Roessler, and Nissenbaum (2019a: 27) point out, manipulation is a ‘success concept.’ A manipulator may have a manipulative intent, but a person is only actually manipulated if the manipulator successfully executed the intended manipulation. This, however, introduces a practical challenge because the health apps that I analyze have tens of millions of users, which results in a sheer endless amount of separate user-app interactions. It is undoable to perform a ‘manipulation check’ for all those separate interactions.

In line with Susser, Roessler, and Nissenbaum (2019a) a practical solution can be offered to this practical problem. Given the scale at which these health apps are deployed (they have tens of millions of users), and given the time and energy they invest in intentionally designing the digital environments they provide to users, I hold that we can be reasonably certain that at least a subset of the users will be successfully manipulated if the manipulative intent is really there. In what follows I will thus focus on *features of the digital environments* provided by health apps to analyze whether those features could be interpreted as betraying a manipulative intent.

4. **For-Profit Health Apps: Between Empowerment and Manipulation?**

In this section I will provide an analysis of the manipulative potential of the digital environments provided by for-profit health apps. Before I proceed, it bears emphasis that the four elements of manipulation I identified *all* need to be present to be able to speak of manipulation. At least two elements of my conception of manipulation – namely manipulation as the *intentional* attempt to make others serve one’s own interests while disregarding or encapsulating the interests of one’s targets – are typical of *any* form of business-to-consumer market interaction in a capitalist society. Companies always intentionally try to further their own ends by ‘making use’ of consumers. Now, because all commercial interactions in the capitalist marketplace share these two elements of manipulation, my argument may be *misconstrued* as stating that 1) all commercial interactions are manipulative, and 2) all user-app interactions are at heart commercial

interactions, so 3) therefore for-profit health apps are problematic because they are commercial in nature. This is why it is so important to emphasize that all four elements of manipulation must be present for a digital environment to qualify as being manipulative. Regular business-to-consumer interactions only rise to the level of manipulation when there is also an intentional attempt to identify and target exploitable characteristics and, moreover, when there is an attempt to obfuscate the intention to manipulate and the means used to do so.³⁶ For my analysis of popular for-profit health apps, this means that the commercial nature of the apps does not, by itself, make them manipulative and therefore problematic; it is only when their commercial nature gives rise to particular commercial practices and app designs that we can possibly speak of manipulation.

Now, in order to understand the potential for manipulation in health apps, we need to acknowledge that people typically use a health app more than one time and sometimes even for an extended period of time. We thus need to understand how the intentional design of the digital environment within which users interact with the app can influence the users *over time*. So instead of analyzing separate, isolated interactions between an app and its users, we need to focus on the way in which the intentionally designed environment is used as a tool to shape the developing relationship between user and app (see Frischmann and Selinger (2018) for a similar argument). This is also why I focus on the potentially manipulative design features of digital environments, rather than on all the

³⁶ It could of course still be argued that quite some business-to-consumer practices rise to the level of manipulation. Take, for example, Santelli (1983) and Crisp (1987) who argue that nearly all advertising is – at least slightly – manipulative. If one really wants to stretch my concept of manipulation, one could even try to argue that a billboard showing advertising is manipulative. Such a billboard with an advertisement for company X is (1) put up intentionally by company X, (2) with the aim to further the ends of company X without a genuine regard for the interests of the people passing by the billboard, (3) is designed by company X in such a manner that it targets either particular people in the street, or particular desires or fears of people in the street, and (4) the billboard does not explicitly communicate that “company X is trying to target you in such a manner that company X’s earns as much as possible.” Even if we agree that a billboard can, strictly speaking, be interpreted to be manipulative, it does not follow that every instance of manipulation warrants the same level of scrutiny. There is a significant difference between, on the one hand, a billboard that displays one and the same message to every person at a fixed location, and, on the other hand, a digital health environment which builds a relationship with the user over time, offers a continuous communication channel to the user, and can be personalized in real time based on what the continuous experiments tell will lead to maximum engagement. Unlike billboards, digital technologies like the health apps discussed can get to know their users over time and can at any time they see fit (e.g. through push notifications) try to leverage that information to manipulate every user personally.

individual, separate interactions between a user and an app that take place within such digital environments.

My main argument is that *an indiscriminate and systematic pursuit* of engagement optimization leads to the use of techniques and designs of digital environments that risk being manipulative and shapes the user-app relationship in an undesirable manner. In practice, however, it remains difficult to draw the line between truly manipulative app designs and unproblematic or even *empowering* app designs. As it turns out, many of the *conditions for* empowering health app designs largely overlap with the conditions for manipulation. For example, a health app that collects and analyzes user data to learn about its users' biases, insecurities, and weaknesses can use that information to personalize the app experience to cater to the specific needs of a user. At the same time, insights into the weaknesses of users can also be *misused* by an app to manipulate them into displaying behavior that solely benefits the app while neglecting or even undermining the interests of the users. The goal of engagement optimization itself can also be part of both empowering and manipulative app experiences. There can be users for whom gentle nudges to continue using a health app can contribute to their health goals. At the same time, however, not *all* users will benefit from the attempts to maximize time spent in a health app. Put simply, the line between empowerment and manipulation can be rather vague in the health app context. In what follows, I will first discuss the separate elements of manipulation in relation to for-profit health apps (sections 4.1.-4.4.) I will then, in section 4.5., briefly reflect on the conditions for manipulation in for-profit health apps and the vague but important line between manipulation and empowerment.

4.1. Targeting Exploitable Characteristics

As I discussed earlier, engagement can be engineered. To do so successfully, health apps – like regular apps – need to know who their users are and what they respond to. Insights into who one's users are and how they interact with one's app inform the (real-time) personalization of the digital environment provided by the app (Yeung 2017). As we saw before, even an ostensibly simple and unobtrusive app like Headspace is – when looking under the hood – run like a never-ending experiment aimed at becoming increasingly

better at identifying which users need to be targeted in what manner to have them spend as much time (and money) as possible on the app. The data analysts need to identify the “key levers for new user growth, retention, and revenue,” and “[l]everage data to understand the product, identify opportunities, and execute initiatives to drive growth and engagement,” and “[d]rive experiment design, interpretation, and actionable insights.”³⁷ What health apps such as Headspace – like other apps – strive for is to *automate* the process of finding out which buttons to push and which strings to pull. Their ability to continuously run large experiments effectively comes down to pushing all known buttons and pulling all known strings and feeding the effects back into their efforts to personalize the app experience in real-time. Put simply, the pursuit of engagement optimization by health apps naturally leads to a situation where identifying and targeting exploitable characteristics of users becomes an essential part of the health apps’ business operations.

Beyond these advanced tools for data analytics which are employed to identify and target exploitable characteristics, the health discourse found in health apps should also be taken into account. As I explained earlier, it is no accident that apps are very eager to frame their goals and content in terms of health. People desire health and people’s desire for health is reinforced by the contemporary culture of healthism which tells them that they *should* be preoccupied with their health (Crawford 2006, Lupton 2013, Cederström & Spicer 2015). Health apps ride this wave by adopting and reproducing a health discourse that emphasizes how important it is to be preoccupied with one’s health and how nearly everything can be considered important to one’s health. Through their imagery and written content health apps try to tap into this “extreme desire for health and wellness” as the advertising agency *NativeAdBuzz* puts it.³⁸ The user’s already existing *general* desire for health is thus targeted by health apps with their own very *particular* health-as-wellness discourses which communicates that nearly everything in a person’s life is relevant to one’s health and can – and should – be managed through an app. Through their health discourse, health apps try to shape their users’ health-identity in such a manner that the users become

³⁷ Headspace, ‘Senior Data Analyst,’ job description that has since been removed, screenshot available here: <https://imgur.com/a/qtSe4li>.

³⁸ NativeAdBuzz, ‘This Health and Wellness Boom Has Been Building for Years... And It’s Finally About to ERUPT (Urgent: Your Free VIP Christmas Gift Has Arrived)’: <http://www.nativeadbuzz.com/blog/this-health-and-wellness-boom-has-been-building-for-years-and-its-finally-about-to-erupt-urgent-your-free-vip-christmas-gift-has-arrived/>

more receptive to the engagement optimization practices. Put differently, one could say that through the deployment of their health discourses, popular commercial health apps attempt to *create* exploitable characteristics in their users which can then, next, be targeted to optimize conversion.³⁹

4.2. *Whose Interests Are Being Served?*

So to the extent that health apps are organized primarily around the principle of engagement optimization, there will also be corresponding attempt to discover and target exploitable characteristics. The next important question from a manipulation perspective is whether popular for-profit health apps try to further their own ends, while either displaying a structural disregard for the true interests of the users or by encapsulating their interests? Because the popular health apps I discuss are commercial enterprises many of which are funded by venture capital,⁴⁰ it is unsurprising that their master aim is to generate profits. The users are a necessary means towards this self-serving end. But do health apps simultaneously display a disregard for the users' true interests? This argument may seem difficult to make because health apps do promise to help users adopt a healthier lifestyle which is generally speaking in the users' interest. I do not want to engage with the empirical question of whether popular for-profit health apps *actually* help people live healthier. Rather, I would like to point out that to the extent that people do in fact live healthier lives as a result of using a health app, this is ultimately a contingent by-product. Above all, most popular for-profit health apps try to get users to engage with the app as much as possible. There is a subtle but important difference between, on the one hand, aiming to maximize engagement (to maximize conversion) *in a systematic and indiscriminate fashion*, and, on the other hand, aiming for optimal health outcomes. If one

³⁹ Already in 1999, Hanson and Kysar used the concept of 'market manipulation' to identify such cases, a concept that was later updated by Calo (2014) who spoke of 'digital market manipulation.' Calo (2014: 1018) noted how "firms will increasingly be able to *create* suckers, rather than waiting for one to be born." Spencer (2019: 34) has argued in a similar vein that "[r]ather than discovering existing vulnerabilities, marketers could exacerbate or even create vulnerabilities in individual subjects and then exploit those vulnerabilities."

⁴⁰ For example, Headspace was funded through four funding rounds, raising \$75 million (<https://www.crunchbase.com/organization/headspace#section-investors>). MyFitnessPal also received funding from venture capitalists (<https://www.crunchbase.com/organization/myfitnesspal#section-investors>) and was later acquired by Under Armour for \$475 million (Olson 2015). Fitbit also saw four funding rounds raising \$66 million from venture capitalists (<https://www.crunchbase.com/organization/fitbit#section-investors>).

optimizes for engagement, it is very well possible that all things considered a significant amount of users will find the health app in question to be helpful. At the same time, however, it is obvious that optimizing for engagement is at best a *very imprecise* proxy for optimizing health outcomes. Consider, for instance, MyFitnessPal and users with (a history of) eating disorders. Trying to get users to engage with MyFitnessPal as much possible can be highly problematic for this group (Eikey & Reddy 2017, Eikey *et al.* 2017). The same holds for people without an eating disorder: maximizing engagement with an app such as MyFitnessPal does *not* necessarily lead to healthier behavior or a happier life, quite to the contrary. Even if one is considered ‘healthy,’ the use of popular for-profit health apps – which often feature idealized discourses on health – can spark “frustration, disappointment, the fear of becoming too controlled, and annoyance or guilt evoked by the demands of the app” (Lupton 2018: 1).⁴¹ These examples shows that optimizing for engagement in a systematic and indiscriminate fashion does not necessarily coincide with furthering the true interests of the users: there is no one-on-one causal relation between spending more time with (and money on) a health app and leading a healthier, happier life. Apps like MyFitnessPal and Headspace thus cleverly *encapsulate* the users’ general interest in being healthy in order to maximize – in an indiscriminate fashion – the engagement of *all* its users with its app, which is not the same as *actually aiming to further* the interests of all their users. So despite the fact that spending time with a health app may appear to be beneficial to its users, it is still accurate to conclude that for-profit health apps display a disregard for the true interests of the users of their app to the extent that they optimize for engagement in an indiscriminate fashion. If a health app’s real and primary aim is to support its users’ health, it should not aim to optimize engagement, since maximum engagement will only be *accidentally* beneficial to *some* of its users. In sum, even though popular for-profit health apps ostensibly seek to further the true interests of their users because everyone wants to be healthy and health apps claim to support people’s healthy

⁴¹ Consider also Culbert, Racine, and Klump’s (2015) meta-analysis of what causes people’s problematic relation to food. They emphasize that especially for perfectly healthy adolescent and young adult females, (digital) media exposure, and more specifically health ideals portrayed in those media, “have all been shown to prospectively predict increased levels of disordered eating cognitions and behaviors (e.g., body dissatisfaction, dieting, bulimic symptoms)” (Culbert, Racine & Klump 2015: 1145).

lifestyles, we should conclude that apps that are run as engagement optimization systems do not in fact seek to further the true health interests of their users.

4.3. *Do Health Apps Attempt to Keep Their Subtle Influences Hidden?*

How about the manipulation requirement that the attempt to influence the decision-making of users is not announced or emphasized by the health app? Here my argument about the deliberate (re)production of a particular health discourse is again important. Popular for-profit health apps adopt a cheerful and positive health-as-wellness discourse because it helps obfuscate to users both the commercial nature of the relationship between the app and the user, and the commercial nature of (some of) the content in the app. I have already argued elsewhere that popular for-profit health apps seek to build ongoing *commercial* relationships with their users, while *de-emphasizing* the commercial nature of those relationships (Sax, Helberger & Bol 2018). Take, for instance, MyFitnessPal and the earlier mentioned MyFitnessPal Blog which is the dominant feature of the app because it takes up the majority of the interface. The blog is rampant with native advertising which is dressed up as genuine health content⁴² and very hard to detect as being advertising. The use of native advertising is unsurprising since professionals in the field of commercial communication and UX design know that people are notoriously bad at recognizing native advertising: “Nearly all the published research on online sponsored content to date clearly shows that a majority of consumers do not realize they are looking at an ad.” (Wojdyski *et al.* 2017: 150) Headspace is another good example. The cheerful, happy design and layout of the app, as well as the texts in the app and on the website, are all aimed at conveying the message that Headspace is just a simple, basic health app containing some mindfulness and meditation exercises. The app is designed *not* to feel like a digital environment that is designed as an advanced engagement optimization system.

⁴² Here are a few examples. The blog post called ‘A Day in the Life of a Yoga Teacher’ (<https://blog.myfitnesspal.com/day-life-yoga-teacher/>) is, in reality, native advertising for skincare products that are framed as being part of a healthy, mindful life. The blog post called ‘These Playlists Were Built to Make You Better’ (<https://blog.myfitnesspal.com/these-playlists-were-built-to-make-you-better/>) is native advertising for a particular brand of headphones. The blog post called ‘Why and How You Should Nix an Alarm Clock’ (<https://blog.myfitnesspal.com/why-and-how-you-should-nix-an-alarm-clock/>) is native advertising for a company offering “certified sleep coaches” and for a company selling a wide range of sleep products.

4.4. *Intentionality*

The last element left to be discussed concerns intentionality: are the (structural elements of the) digital environments provided by health apps, and the effects they can have on users, the result of intentional choices on the part of health app providers? The claim I seek to defend here is that many popular for-profit health apps can be considered manipulative digital *environments*, without it necessarily being the case that *all* of the literally billions of separate interactions between users and health apps are – when analyzed in isolation – manipulative. The design of the environments is, as far as I can judge, a completely intentional process. Based on, for instance, the job descriptions that are available for all popular commercial health apps, it is clear that even the health apps that are considered to be very simple, basic, and benign – such as Headspace – are actively engineering and tweaking the entire digital environment they control to optimize conversion.

4.5. *The Manipulative Potential of Popular For-Profit Health Apps*

Popular for-profit health apps try to build ongoing commercial relationships with as many users as possible and employ a range of techniques to build digital environments that help them to optimize user retention, engagement, and conversion. In this regard, health apps do not differ from other types of popular apps, such as the gaming apps Candy Crush or Pokémon Go. Such gaming apps, like popular for-profit health apps, are also built and run as optimization systems. There is, of course, one striking *difference* between health apps and popular gaming apps like Candy Crush: health apps deal with *health*, for which people have an inescapable need and universal desire. When a value as important as health is used for commercial gain, we should be especially vigilant, for *use* of such a value can easily turn into *misuse*. Put differently, it is precisely *because* popular for-profit health apps are built and run like any other type of for-profit app (such as gaming apps) that we should be vigilant.

My analysis of the four different requirements for manipulation provided above does indeed show that there exists a serious manipulative potential in the digital environments provided by popular for-profit health apps. In practice, however, it can be difficult to clearly separate for-profit health apps' manipulative potential from their

potential to help users. This is due to the fact that the conditions for support or empowerment largely overlap with the conditions for manipulation. The data-driven nature of health apps, as well as their ability to run experiments and dynamically change (parts of) the interface or functionalities, allow health apps to identify and target exploitable characteristics of their users to make them serve the health app providers' ends. This is where the potential for manipulation clearly shows itself. But those very same technological capabilities can also be used to identify and target characteristics of their users for the purpose of personalized interventions and suggestions aimed at helping the user. So to identify and evaluate the manipulative potential of health apps, it is not enough to focus on the technological affordances of health apps as such. It is equally important to look at how those technological affordances are put to use.

I would like to suggest that in the health app context, we should mainly focus on the extent to which commercially-inspired optimization logics have been implemented. When a for-profit health app is organized and run as an optimization system that seeks to *systematically and indiscriminately* optimize for engagement, retention, and conversion, the potential for empowerment will quickly turn into a potential for manipulation. When the pursuit of optimization of engagement, retention, and conversion is systematic and indiscriminate, users are treated as instruments whose interests can be encapsulated to help further the business ends of the app providers. In such cases, the goal of engagement optimization tends to lead to the development and deployment of data-driven systems that (attempt to) automate the process of identifying and characteristics and circumstances of their users *for the purpose of exploiting those characteristics and circumstances for commercial gain*. Although users of such apps may – as an ultimately contingent by-product – become healthier in the process of using an app, this is not the real aim of the providers of popular for-profit health apps. Health apps, moreover, know all too well that people desire health and that this desire can be used to obfuscate the commercial nature of the user-app relationship.

So when health apps optimize for engagement systematically and indiscriminately, they knowingly introduce and accept the risk that they will exert manipulative influences on their users. It also follows that for-profit health apps' pursuit of profit and the use of clever technologically driven techniques, are not inherently problematic or manipulative. A

health app that collects and analyzes user data to build user profiles can also use those profiles to decide which users *not* to target with nudges aimed at maximizing engagement and retention. Big Tech companies (including popular for-profit health apps) like to brag about their ability to ‘get to know’ their users to anticipate their behavior, needs, and desires. In the health app context, where the value of health is at stake, we may want to require for-profit health app to observe (something akin to) a duty of care: if you collect large amounts of user data, you are also responsible for *not* targeting health-related weaknesses in order to dispose those users to your commercial ends. As I have argued elsewhere, European consumer law can potentially be used to regulate the manipulative potential of health apps by, for instance, imposing (something akin to) a duty of care on health app providers, or by banning indiscriminate manipulative data-driven targeting practices (Sax, Helberger & Bol 2018).

5. Conclusion

Recently there has been an increasing amount of attention for the engineering of highly persuasive digital environments (e.g. Yeung 2017, Selinger & Frischmann 2018, Zuboff 2018). Some have argued that such digital environments foster a unique manipulative potential (e.g. Susser, Roessler & Nissenbaum 2019a, 2019b). In this article I discussed a subset of health apps – namely popular for-profit health apps – as a paradigm case of digital environments that risk being manipulative.

The risk of manipulation can be explained by a combination of multiple factors. First, health is essential to and desired by everyone, which means that appeals to health can be *misused* to subtly influence behavior of health app users in ways they – all things considered – may not appreciate. Second, popular for-profit health apps often appear to optimize for user engagement in order to maximize conversion. Third, to optimize engagement popular for-profit health try to figure out how to best target their users and (dynamically) adept the digital environment. Fourth, most popular for-profit health apps propagate a health discourse that intends to mask the profit-driven architecture of the digital health environment to the users by emphasizing how the app is there to empower a user in her pursuit of health and wellness. Such digital health environments display a lack

of respect for the autonomy of their users by treating them as pawns that can be used by the health app provider, for the benefit of the health app provider, while appealing to a universally necessary good – health – to obfuscate to users the true nature of their (ongoing) *commercial* relationship with the app. Precisely because health is so important to every person, we should be especially vigilant when appeals to health are used to subtly dispose health app users to the self-serving ends of providers of popular for-profit health apps. The fact that a part of the health app user population finds the health apps they use (somewhat) helpful should not lead us to ignore the manipulative designs of some of these apps. Health apps should not be designed in the same way as for instance popular gaming apps are designed, with a hyperfocus on getting as many users as possible ‘hooked’ on the app in order to maximize their engagement with profit-generating features and content. Health is too important to be misused for cynical commercial gains pursued under the shroud of empowerment.

It bears emphasis that my analysis does not imply that all health apps are manipulative, or that all health apps risk being manipulative. It is only *to the extent* that a particular app features all the elements of manipulation discussed that a case can be made for the presence problematic manipulative tendencies. So far from being a critique of all health apps, this article serves as a critique on a *particular subset of health apps with particular characteristics*. It is my hope that my analysis helps in addressing the problems of this particular subset of health apps so that all those health apps that are not manipulative can continue to empower users in a respectful manner.

Acknowledgements

This article was written during a research visit to The Digital Life Initiative at Cornell Tech. I would like to thank Helen Nissenbaum, Eran Toch, Elizabeth O’Neill, Jake Goldenfein, Michael Byrne, Erica Du, Yvonne Wang, Lauren van Haaften-Schick, Mason Marks, and Jessie Taft for providing me with a welcoming and intellectually stimulating environment during my research visit. Special thanks to Elizabeth O’Neill, Eva Groen-Reijman, Thomas Nys, and Beate Roessler for extensive feedback on earlier versions of this article. I would also like to thank the anonymous reviewers for their constructive reviews and helpful

suggestions for improvement of the article.

References

- Ajana, B. (2017). Digital Health and the Biopolitics of the Quantified Self. *Digital Health*, 3, 1-18.
- Alter, A. (2017). *Irresistible: The Rise of Addictive Technology and the Business of Keeping Us Hooked*. New York: Penguin Press.
- Anderson, J. H. (2014). Autonomy and Vulnerability Entwined. In C. Mackenzie, W. Rogers & S. Dodds (Eds.), *Vulnerability: New Essays in Ethics and Feminist Philosophy* (pp. 134-161). Oxford: Oxford University Press.
- Anderson, J. H., & A. Honneth (2005). Autonomy, Vulnerability, Recognition, and Justice. In J. Christman & A. Honneth (Eds.), *Autonomy and the Challenges to Liberalism* (pp. 127-149). Cambridge University Press.
- Armstrong, D. (1995). The Rise of Surveillance Medicine. *Sociology of Health & Illness*, 17(3), 393-404.
- Barnhill, A. (2014). What Is Manipulation? In C. Coons & M. Weber (Eds.), *Manipulation: Theory and Practice* (pp. 51-72). Oxford: Oxford University Press.
- Baron, M. (2003). Manipulativeness. *Proceedings and Addresses of the American Philosophical Association*, 77(2), 37-54.
- Faden, R. R., & Beauchamp, T. L. (1986). *A History and Theory of Informed Consent*. New York: Oxford University Press.
- Boorse, C. (1975). On the Distinction between Disease and Illness. *Philosophy and Public Affairs*, 5(1), 49-68.
- Boorse, C. (1977). Health as a Theoretical Concept. *Philosophy of Science*, 44(4), 542-573.
- Bovens, L. (2009). The Ethics of Nudge. In T. Grüne-Yanoff & S. Hansson (Eds.), *Preference Change: Approaches from Philosophy, Economics and Psychology*. Springer.
- Brodesser-Akner, T. (2018). How Goop's Haters Made Gwyneth Paltrow's Company Worth \$250 million: Inside the Growth of the Most Controversial Brand in the Wellness Industry. *The New York Times Magazine*, July 25, 2018. Available at: <https://www.nytimes.com/2018/07/25/magazine/big-business-gwyneth-paltrow-wellness.html>
- Buss, S. (2005). Valuing Autonomy and Respecting Persons: Manipulation, Seduction, and the Basis of Moral Constraints. *Ethics*, 115(2), 195-235.
- Calo, R. (2014). Digital Market Manipulation. *George Washington Law Review*, 82(4), 995-1051.
- Casper, M., & Morrison, D. (2010). Medical Sociology and Technology: Critical Engagements. *Journal of Health and Social Behavior*, 51(1), 12-32.
- Cederström, C. & Spicer, A. (2015). *The Wellness Syndrome*. Cambridge: Polity Press.
- Chaykowski, K. (2017). Meet Headspace, the App that Made Meditation a \$250 million Business. *Forbes*, January 8, 2017. Available at <https://www.forbes.com/sites/kathleenchaykowski/2017/01/08/meet-headspace-the-app-that-made-meditation-a-250-million-business/#7641f4f81f1b>.
- Christman, J. (1991). Autonomy and Personal History. *Canadian Journal of Philosophy*, 21(1), 1-24.

- Christman, J. (2004). Relational Autonomy, Liberal Individualism, and the Social Constitution of Selves. *Philosophical Studies*, 117(1/2), 143-164.
- Code, L. (1991). *What Can She Know? Feminist Theory and the Construction of Knowledge*. Ithaca: Cornell University Press.
- Cohen, S. (2018). Manipulation and Deception. *Australian Journal of Philosophy*, 96(3), 483-497.
- Crawford, R. (2006). Health as a Meaningful Social Practice. *Health*, 10(4), 401-420.
- Crisp, R. (1987). Persuasive Advertising, Autonomy, and the Creation of Desire. *Journal of Business Ethics*, 6(5), 413-418.
- Culbert, K. M., Racine, S. E., & Klump, K. L. (2015). Research Review: What We Have Learned about the Causes of Eating Disorders – A Synthesis of Sociocultural, Psychological, and Biological Research. *Journal of Child Psychology and Psychiatry* 56(11), 1141-1164.
- Eikey, E. V., & Reddy, M. C. (2017). “It’s Definitely Been a Journey”: A Qualitative Study on How Women with Eating Disorders Use Weight Loss Apps. *Proceedings on the 2017 CHI Conference on Human Factors in Computing Systems*, 642-654.
- Eikey, E.V., Reddy, M.C., Booth, K.M., Kvasny, L., Blair, J.L., Li, V., & Poole, E. (2017). Desire to be Underweight: Exploratory Study on Weight Loss App Community and User Perceptions of the Impact on Disordered Eating Behaviors. *JMIR mHealth and uHealth*, 5(10), e150.
- European Commission (2012). *Communication from the Commission to the Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions, Health Action Plan 2012–2020 – Innovative healthcare for the 21st century (Brussels, 6.12.2012. COM/2012/0736 final)*. Brussels: European Union.
- European Commission (2014). *Green paper on mobile health* (Brussels, 10.4.2014, COM(2014) 135 final). Available at <https://ec.europa.eu/digital-single-market/en/news/green-paper-mobile-health-mhealth>
- Eyal, N. (2014). *Hooked: How to Build Habit-Forming Products*. New York: Portfolio/Penguin.
- Fahy, R., Van Hoboken, J., & Van Eijk, N. (2018). Data Privacy, Transparency and the Data-Driven Transformation of Games to Services. *IEEE Games, Entertainment, Media Conference*, 1-9.
- Floridi, L. (2005). The Ontological Interpretation of Informational Privacy. *Ethics and Information Technology*, 7(4), 185-200.
- Fotopoulou, A., & O’Roirdan, K. (2017). Training to Self-Care: Fitness Tracking, Biopedagogy and the Healthy Consumer. *Health Sociology Review*, 26(1), 54-68.
- Foucault, M. (1975). *The Birth of the Clinic: An Archeology of Medical Perception*. New York: Vintage Books.
- Frischmann, B., & Selinger, E. (2018). *Re-engineering Humanity*. Cambridge: Cambridge University Press.
- Gorin, M. (2014). Towards a Theory of Interpersonal Manipulation. In C. Coons & M. Weber (Eds.), *Manipulation: Theory and Practice* (pp. 73-97). Oxford: Oxford University Press.
- Greenspan, P. (2003). The Problem with Manipulation. *American Philosophical Quarterly*, 40(2), 155-164.

- Hanson, J. D., & Kysar, D. A. (1999). Taking Behavioralism Seriously: The Problem of Market Manipulation. *New York University Law Review*, 74(3), 630-749.
- Hardin, R. (2002). *Trust and Trustworthiness*. New York: Russell Sage Foundation.
- Hayek, F. A. (2006 [1960]). Freedom and Coercion. In D. L. Miller (Ed.), *The Liberty Reader* (pp. 80-99). New York: Routledge.
- Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.
- Lanzing, M. (2016). The Transparent Self. *Ethics and Information Technology*, 18(1), 9-16.
- Lanzing, M. (2018). "Strongly Recommended": Revisiting Decisional Privacy to Judge Hypernudging in Self-Tracking Technologies. *Philosophy & Technology*, online first.
- Lupton, D. (2012). M-Health and Health Promotion: The Digital Cyborg and Surveillance Society. *Social Theory & Health*, 10(3), 229-244.
- Lupton, D. (2013). Quantifying the Body: Monitoring and Measuring Health in the Age of mHealth Technologies. *Critical Public Health*, 23(4), 393-403.
- Lupton, D. (2018). I Just Want It to Be Done, Done, Done! Food Tracking Apps, Affects, and Agential Capacities. *Multimodal Technologies Interact*, 2(2), 29-44.
- Mackenzie, C. & N. Stoljar (2000). 'Introduction: Autonomy Refigured' pp. 3-31 in C. Mackenzie & N. Stoljar (eds.) *Relational Autonomy: Feminist Perspectives on Autonomy, Agency, and the Social Self*. Oxford: Oxford University Press.
- Mayes, C. (2015). *The Biopolitics of Lifestyle: Foucault, Ethics and Healthy Choices*. London: Routledge.
- Meyers, D. (1989). *Self, Society and Personal Choice*. Oxford: Oxford University Press.
- Mills, C. (2014). Manipulation as an Aesthetic Flaw. In C. Coons & M. Weber (Eds.), *Manipulation: Theory and Practice* (pp. 135-150). Oxford: Oxford University Press.
- Nedelsky, J. (1989). Reconceiving Autonomy: Sources, Thoughts and Possibilities. *Yale Journal of Law and Feminism*, 1, 7-36.
- Noggle, R. (1996). Manipulative Actions: A Conceptual and Moral Analysis. *American Philosophical Quarterly*, 33(1), 43-55.
- Nordenfelt, L. (1986). Health and Disease: Two Philosophical Perspectives. *Journal of Epidemiology and Community Health*, 41, 281-284.
- Nordenfelt, L. (1987) *On the Nature of Health: An Action Theoretic Approach*. Dordrecht: Reidel.
- Nys, T. R. V. (2016). Autonomy, Trust, and Respect. *Journal of Medicine and Philosophy*, 41(1), 10-24.
- Olson, P. (2014). MyFitnessPal Starts Tracking Steps to Grow the World's Largest Nutrition Database. *Forbes*, May 1, 2014. Available at <https://www.forbes.com/sites/parmyolson/2014/05/01/myfitnesspal-starts-tracking-steps-to-grow-the-worlds-largest-nutrition-database/#341b09d05968>.
- Olson, P. (2015). Under Armour Buys Health-Tracking App MyFitnessPal for \$475 Million. *Forbes*, February 4, 2015. Available at <https://www.forbes.com/sites/parmyolson/2015/02/04/myfitnesspal-acquisition-under-armour/#1a75350c6935>.
- Overdorf, R., Kulnych, B., Balsa, E., Troncoso, C., & Gürses, S. (2018). POTs: Protective Optimization Technologies. *arXiv: 1806.02711*.
- Patterson, H. (2013). Contextual Expectations of Privacy in Self-Generated Health Information Flows. *The 41st Research Conference on Communication, Information and Internet Policy*. Available at <https://ssrn.com/abstract=2242144>.

- Petersen, A., & Bunton, R. (Eds.) (1997). *Foucault, Health, and Medicine*. London: Routledge.
- Rudinow, J. (1978). Manipulation. *Ethics*, 88(4), 338-347.
- Sanders, R. (2017). Self-Tracking in the Digital Era: Biopower, Patriarchy, and the New Biometric Body Projects. *Body & Society*, 23(1), 36-63.
- Santilli, P. C. (1983). The Informative and Persuasive Functions of Advertising: A Moral Appraisal. *Journal of Business Ethics*, 2(1), 27-33.
- Sax, M., Helberger, N., & Bol, N. (2018). Health as a Means Towards Profitable Ends: mHealth Apps, User Autonomy, and Unfair Commercial Practices. *Journal of Consumer Policy*, 41(2), 103-134.
- Spencer, S. B. (2019). The Problem of Online Manipulation. Available at <https://ssrn.com/abstract=3341653>
- Stoljar, N. (2011). Informed Consent and Relational Conceptions of Autonomy. *The Journal of Medicine and Philosophy*, 36(4), 375-384.
- Susser, D. (2019). Invisible Influence: Artificial Intelligence and the Ethics of Adaptive Choice Architectures. *AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*, January 27-28, Honolulu, Hawaii, USA. Available at: http://www.aies-conference.com/wp-content/papers/main/AIES-19_paper_54.pdf
- Susser, D., Roessler, B., & Nissenbaum, H. (2019a). Online Manipulation: Hidden Influences in a Digital World. *Georgetown Law Technology Review*, 4(1), 1-45.
- Susser, D., Roessler, B., & Nissenbaum, H. (2019b). Technology, Autonomy, and Manipulation. *Internet Policy Review*, 8(2).
- Thaler, R. H. & Sunstein, C. R. (2008). *Nudge: Improving Decision About Health, Wealth, and Happiness*. New Haven, CT: Yale University Press.
- Wojdyski, B. W., & Evans, N. J. (2016). Going Native: Effects of Disclosure Position and Language on the Recognition and Evaluation of Online Advertising. *Journal of Advertising*, 45(2), 157-168.
- Wojdyski, B. W., Bang, H., Keib, K., Jefferson, B. N., Choi, D., & Malson, J. L. (2017). Building a Better Native Advertising Disclosure. *Journal of Interactive Advertising*, 17(2), 150-161.
- Wood, A. W. (2014). Coercion, Manipulation, Exploitation. In C. Coons & M. Weber (Eds.), *Manipulation: Theory and Practice* (pp. 17-50). Oxford: Oxford University Press.
- World Health Organization (2011). *mHealth. New horizons for health through mobile technologies. Report based on the findings of the second global survey on eHealth*. Geneva: World Health Organization.
- Yeung, K. (2017). 'Hypernudge': Big Data as a Mode of Regulation by Design. *Information, Communication & Society*, 20(1), 118-136.
- Zuboff, S. (2015). Big Other: Surveillance Capitalism and the Prospect of an Information Civilization. *Journal of Information Technology*, 30(1), 75-89.
- Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the Frontier of Power*. New York: Public Affairs.